

Dell EMC Data Domain Boost for Partner Integration

Version 3.4

Administration Guide

302-003-982

REV. 01

Copyright © 2017 Dell Inc. and its subsidiaries. All rights reserved.

Published May 2017

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

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

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

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

CONTENTS

	Preface	5
Chapter 1	Introducing Data Domain Boost	9
	Revision History.....	10
	Overview of DD Boost.....	10
	Supported Configurations	11
	Upgrade Compatibility.....	11
Chapter 2	DD Boost Features	13
	Overview of DD Boost Features.....	14
	Distributed Segment Processing	14
	In-flight Encryption.....	15
	Global Authentication and Encryption.....	16
	Methods of Setting Authentication and Encryption.....	16
	Authentication and Encryption Settings.....	16
	Authentication and Encryption Options.....	17
	Backwards Compatibility Scenarios.....	19
	Managed File Replication (MFR).....	21
	Low-Bandwidth Optimization.....	21
	Encrypted Managed File Replication.....	21
	DD Boost and High Availability.....	22
	DD Boost, HA, and failover.....	22
	Partial HA configurations	23
	MTree Replication.....	23
	IPv6 Support.....	23
	Dynamic Interface Groups: DD Boost IP Data Path Management.....	24
	Interfaces.....	25
	Clients.....	27
	Using interface groups for Managed File Replication (MFR).....	29
	IP Failover Hostname.....	30
	DD Boost-over-Fibre Channel Transport.....	31
	DD Boost-over-Fibre Channel Path Management.....	33
	Initial Path Selection.....	35
	Dynamic Re-Balancing.....	35
	Client Path Failover.....	35
	Queue-Depth Constraints.....	36
	Virtual Synthetic Backups.....	36
	Client Access Validation.....	37
	DD Boost Multiuser Data Path.....	37
	Storage Unit Management.....	37
	Multiuser Storage Units Access Control.....	37
	Storage Unit Capacity Quotas.....	38
	Storage Units Stream Count Management.....	38
Chapter 3	Preparing the Data Domain System for DD Boost	41
	Enabling DD Boost on a Data Domain System.....	42
	Assigning Multiple Users to DD Boost.....	42
	Creating Storage Units	43

	Configuring Logical Quotas for Storage Units (Optional)	44
	Configuring Storage Units with Stream Limits (Optional).....	45
	Configuring Distributed Segment Processing.....	46
	Configuring Dynamic Interface Groups	47
	Modifying an Interface Group.....	49
	Removing an Interface Group.....	50
	Using Dynamic Interface Groups for MFR.....	50
	Replication over LANs.....	51
	Replication over WANs.....	53
	Other Supported Use Cases.....	54
	Replication Failover Hostname.....	57
	Network Address Translation (NAT) Support.....	58
	Resolving Backup/Replication Conflicts.....	59
	Configuring MFR.....	60
	Enabling Low-Bandwidth Optimization	60
	Enabling Encryption.....	60
	Enabling IPv6 Support.....	61
	Changing the MFR TCP Port.....	61
	Configuring Client Access Validation.....	61
	Configuring DD Boost-over-FC Service.....	62
	Sizing DD Boost-over FC Device-Set.....	64
	Sizing Calculation.....	65
	Installing the AIX DDdfc Device Driver (Optional for AIX Clients)....	67
	Configuring the SCSI Generic Device Driver for Solaris Clients.....	68
	Setting Global Authentication and Encryption.....	70
	Showing Global Authentication and Encryption Settings.....	70
	Resetting Global Authentication and Encryption Settings.....	70
Chapter 4	Backup Application Administration	71
	Configuring a Backup Server.....	72
	Backup Administration.....	72
	Network Time-Outs.....	72
Chapter 5	Basic Troubleshooting	73
	General Troubleshooting.....	74
	Data Domain System Settings for File Replication.....	74
	Resolve time-out error.....	74
	Managed File Replication Job Fails.....	74
	Add license for Replication.....	74
	Verify Encrypted Managed File Replication Configuration.....	75
	Virtual Synthetic Backup.....	75

Preface

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

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

Note

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

Purpose

This guide explains how to configure and use EMC Data Domain Boost when used with partner applications, including:

- Dell NetVault Backup
 - Dell vRanger Pro
 - EMC Avamar
 - EMC Database application agent for DD Boost for Enterprise Applications and ProtectPoint
 - EMC Microsoft application agent for DD Boost for Enterprise Applications
 - EMC NetWorker
 - Hewlett-Packard HP Data Protector
 - Pivotal Greenplum Data Computing Appliance
 - Quest NetVault
 - Quest vRanger Pro
 - Veeam Backup and Replication
 - VMware vSphere Data Protection Advanced (VDPA)
-

Note

A separate guide, the *EMC Data Domain Boost for OpenStorage Administration Guide*, has been published specifically for use with Symantec backup applications (NetBackup and Backup Exec). Consult that publication for guidance on using DD Boost with Symantec OpenStorage.

Audience

This guide is for system administrators who are familiar with backup applications and general backup administration.

Related EMC documentation

The following Data Domain system documents provide additional information:

- *EMC Data Domain Operating System Release Notes*

- *EMC Data Domain Operating System Initial Configuration Guide*
- *EMC Data Domain Operating System Administration Guide*
- *EMC Data Domain Operating System Command Reference Guide*
- *EMC Data Domain Expansion Shelf Hardware Guide*
- The Data Domain system installation and setup guide for each of the supported platforms (for example DD890, DD690g, and so forth).

Where to get help for EMC products

EMC product documentation

To view documentation, go to the EMC Online Support site at <https://support.emc.com> and click **Support by Product** below the Search box. Type a product name in the **Find a Product** box, wait for those words to appear in the list of matches below the box, and click the words. Then click **»**. In the list of categories under the Search box, click **Documentation**.

- The **Product** choices let you filter results by system model name or number or by an associated software release.
- The **Content Type** choices let you filter results by category. Click **More** under **Content Type** to see all of the categories. The categories that contain end-user and compatibility documentation are these:
 - **Manuals and Guides**, for the software and hardware manuals for your system and for integration guides that explain how to use your systems with backup software and other products.
 - **Release Notes**, for specific versions of the EMC operating system in use with your system.
 - **Compatibility Documents**, for guides that show which EMC and third-party components are compatible.

EMC technical support

Go to EMC Online Support at <https://support.emc.com> and click **Service Center**. You will see several options for contacting EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Where to get help for partner applications

Documentation for backup applications is available through the backup application vendor.

Special notice conventions used in this document

EMC uses the following conventions for special notices:

NOTICE

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

Note

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

Typographical conventions

EMC uses the following type style conventions in this document:

Table 1 Typography

Bold	Indicates interface element names, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Highlights publication titles listed in text
Monospace	Indicates system information, such as: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Highlights a variable name that must be replaced with a variable value
Monospace bold	Indicates text for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections—the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your feedback about this document to: DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Introducing Data Domain Boost

This chapter contains the following topics:

- [Revision History](#)10
- [Overview of DD Boost](#)..... 10
- [Supported Configurations](#) 11
- [Upgrade Compatibility](#)..... 11

Revision History

The following table presents the revision history of this document.

Table 2 Revision History of DD Boost for Partner Integration Release 3.4

Revision	Date	Description
01 (3.4)	May 2017	<p>This revision includes information about these new features and tools:</p> <ul style="list-style-type: none"> • Global authentication and encryption options to better protect against man-in-the-middle (MITM) attacks. • This revision includes the following new APIs: <ul style="list-style-type: none"> ▪ ddp_readlink ▪ ddp_symlink ▪ ddp_lstat ▪ ddp_lstat64

Overview of DD Boost

EMC Data Domain Boost (DD Boost) enables backup servers to communicate with storage systems without the need for Data Domain storage systems to emulate tape. The software has two components:

- DD Boost libraries that you install on each backup server to integrate with the DD Boost server that runs on the Data Domain system.
- The DD Boost server that runs on Data Domain systems.

Note


A Data Domain system can be a single Data Domain system, a gateway, or a DD Extended Retention system. With DD Boost 3.2.1 and DD OS 5.7, it can also be a dual-node Data Domain HA system.

The backup application sets policies that control when backups and duplications occur. Administrators manage backup, duplication, and restores from a single console and can use all of the features of DD Boost, including WAN-efficient replicator software. The application manages all files (collections of data) in the catalog, even those created by the Data Domain system.

The Data Domain system exposes pre-made disk volumes called storage units to a DD Boost-enabled backup server. Multiple backup servers, each with the DD Boost libraries, can use the same storage unit on a Data Domain system as a storage server. Each backup server can run a different operating system, provided that the operating system is supported by Data Domain and the backup application.

The figure shows an example configuration of Data Domain Boost.

Figure 1 DD Boost — Configuration



1. Clients
2. Server
3. Primary Storage
4. Backup Server
5. DD Boost Libraries
6. Data Domain
7. Data Domain Appliance
8. DD Boost
9. WAN
10. Secondary Data Domain Appliance
11. Archive to Tape as Required
12. Backup
13. Retention/Restore
14. Replication
15. Disaster Recovery

Supported Configurations

EMC Data Domain supports DD Boost on all Data Domain systems.

The DD Boost library versions must be compatible with the software version of your Data Domain system and with backup application configurations. Data Domain does not support combinations other than those detailed in the *Data Domain Boost Compatibility Guide* available at the EMC Online Support site <https://support.emc.com>.

Upgrade Compatibility

The Data Domain policy of upgrade compatibility for replication is as follows:

- All maintenance and patch versions within a *family* are backward compatible. A family is identified by the first two digits of the release number, such as 5.2. For example, 5.2.0.0, 5.2.0.2, 5.2.1.0, and 5.2.2.0 are all backward compatible.
- Replication is backward compatible across two consecutive release families, such as 5.5 and 5.4, although only the current release within each family is fully tested.
- Both source and destination Data Domain systems must be licensed for replication.

CHAPTER 2

DD Boost Features

New and enhanced capabilities are available for Single Node and DD Extended Retention. DD High Availability (HA) is also supported.

This chapter describes the major features and functionality of the DD Boost software in the following topics:

- [Overview of DD Boost Features](#)..... 14
- [Distributed Segment Processing](#) 14
- [In-flight Encryption](#)..... 15
- [Global Authentication and Encryption](#)..... 16
- [Managed File Replication \(MFR\)](#)..... 21
- [DD Boost and High Availability](#)..... 22
- [MTree Replication](#)..... 23
- [IPv6 Support](#)..... 23
- [Dynamic Interface Groups: DD Boost IP Data Path Management](#)..... 24
- [IP Failover Hostname](#)..... 30
- [DD Boost-over-Fibre Channel Transport](#)..... 31
- [DD Boost-over-Fibre Channel Path Management](#)..... 33
- [Virtual Synthetic Backups](#)..... 36
- [Client Access Validation](#)..... 37
- [DD Boost Multiuser Data Path](#)..... 37
- [Storage Unit Management](#)..... 37

Overview of DD Boost Features

Backup applications are a critical component of data recovery and disaster preparedness strategies. Each strategy requires a strong, simple, and flexible foundation that enables users to respond quickly and manage operations effectively.

EMC Data Domain systems integrate easily with backup software and provide retention and recovery benefits of inline deduplication. Additionally, Data Domain systems provide replication protection over the WAN for offsite disaster recovery.

DD Boost increases performance by distributing the deduplication process between the client and the backup server.

Note

DD Boost performance can vary depending on the type of hardware on which the DD Boost client is running. Best performance is seen with clients running on Intel x86-based family processors. Due to architectural limitations, poorer performance may be seen on non-x86-based systems such as Itanium (HP-UX), Power (AIX), and Sparc (Solaris).

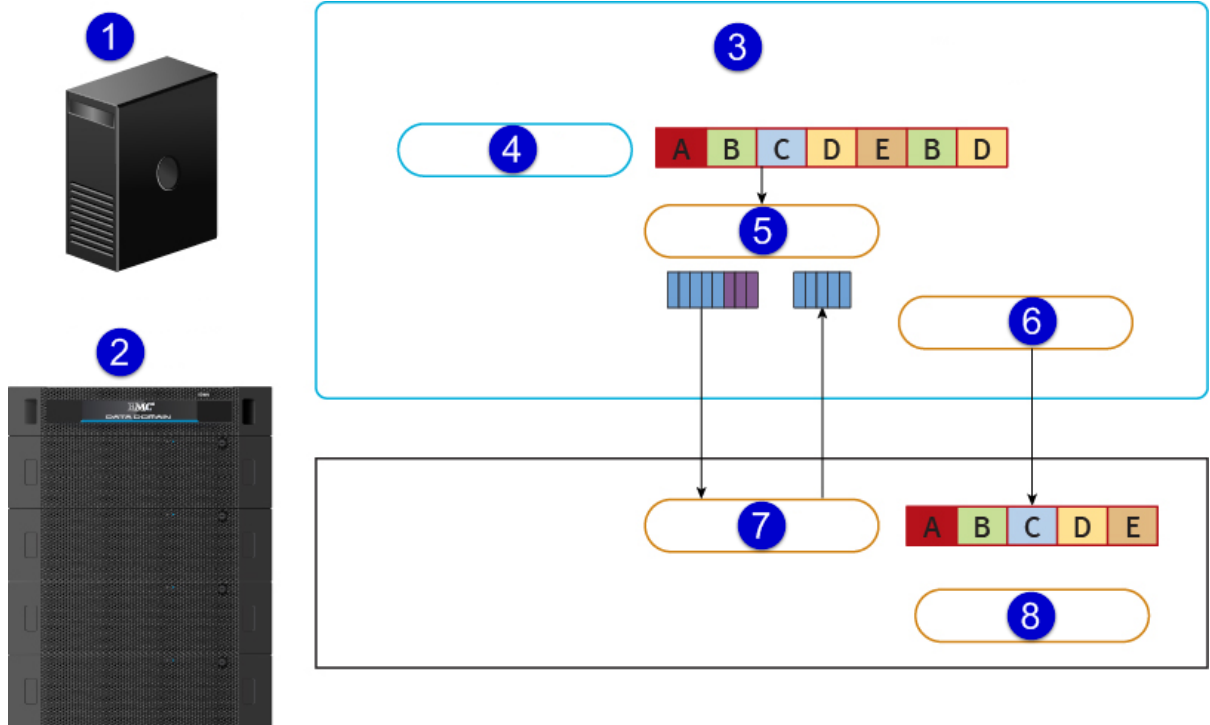
Distributed Segment Processing

The distributed segment processing functionality of the DD Boost software distributes the deduplication process between client and server to avoid sending duplicate data to the Data Domain system.

Distributed segment processing provides the following benefits:

- Potentially lower network traffic generation because the DD Boost Library sends only unique data to a Data Domain system. In general, the greater the redundancy in the data set, the greater the saved network bandwidth to the Data Domain system.
- With distributed segment processing, the DD Boost Library uses 24 MB of memory for every file backed up. DD Boost 5.7 with HA uses 128 MB of memory for every file backed up.

Figure 2 Distributed Segment Processing Enabled



1. Database Server
2. Data Domain System
3. Data Domain Boost libraries
4. Segment
5. Fingerprint
6. Compress
7. Filter
8. Write

Note

When using the Solaris 11/11.1 bundled OpenSSL 1.0.0.j and running on either Solaris 11 (with SRU2 or later) or Solaris 11.1 or later, the plug-in offers improved distributed segment processing (DSP) compared to other Solaris systems. DSP is now enabled by default for Solaris plug-ins running on a SPARC T4 or T5 processor and running Solaris 11 (with SRU2 or later) or Solaris 11.1 or later.

In-flight Encryption

In-flight encryption allows applications to encrypt in-flight backup or restore data over LAN from the Data Domain system. This feature was introduced to offer a more secure data transport capability.

When configured, the client is able to use TLS to encrypt the session between the client and the Data Domain system. The specific cipher suite used is either ADH-AES256-SHA, if the HIGH encryption option is selected, or ADH-AES128-SHA, if the MEDIUM encryption option is selected.

Global Authentication and Encryption

DD Boost offers global authentication and encryption options to defend your system against man-in-the-middle (MITM) attacks.

The global options ensure new clients are protected, but also allow you to configure different values for each client. In addition, client settings can only strengthen security, not reduce it.

Setting the global authentication mode and encryption strength establishes minimum levels of authentication and encryption. All connection attempts by all clients must meet or exceed these levels.

Note

These measures are not enabled by default; you must change the settings manually.

The default global options are backwards-compatible, meaning:

- You do not have to update the DD Boost library. All existing clients and applications will perform in the same manner with the default settings of the new options.
- There is no impact on performance because there is no added encryption.
- Clients and applications that use certificates with transport layer security (TLS) can continue to work with no changes.

Note

If the global settings are different than the default settings, existing clients might need to be updated.

Methods of Setting Authentication and Encryption

You can specify authentication and encryption settings in three ways.

- **Connection request**
You do this by using the `ddp_connect_with_config` API in the client application.
- **Per-client settings**
You do this by using CLI commands on the Data Domain system.
- **Global settings**
You do this by using CLI commands on the Data Domain system.

If both per-client and global values are set, the stronger or higher setting is used.

Authentication and Encryption Settings

You can consider several factors when deciding authentication and encryption settings. However, it is recommended that you always choose the maximum available setting for maximum security.

Maximum security will impact performance. If you have a controlled environment where maximum security is not required, you might want to use other settings.

Global settings

The global setting determines the minimum levels of authentication and encryption. Connection attempts that do not meet these criteria will fail.

Per-client settings

If the setting is defined on a per-client basis, the setting you choose must either match or be greater than the maximum per-client authentication setting and the maximum global authentication setting.

For example:

- If a client is configured to require "two-way password" authentication and the global authentication setting is two-way TLS, then two-way TLS authentication must be used.
- If the client is configured with the authentication setting "two-way TLS" and the global setting is "two-way passwords", then "two-way TLS" must be used.

Caller-specified values

If the caller-specified values are lower than either the global or per-client settings, the connection is not allowed. However, if the caller-specified values are higher than the global or per-client settings, the connection will be made using the caller-specified values.

For example, if the caller specifies "two-way-password" but either the global or per-client value is "two-way," the connection attempt fails. However, if the caller specified "two-way" and the global and per-client values are "two-way-password," "two-way" authentication is used.

Authentication and Encryption Options

You can select one of three allowed settings for both the global and authentication and encryption settings.

For the per-client settings, five authentication settings are allowed and three encryption settings (the same encryption settings as those for global).

Note

Authentication and encryption values must be set at the same time due to dependencies.

Global authentication and encryption options

You have a range of choices with the options `global-authentication-mode` and `global-encryption-strength`.

Authentication settings

The following list ranks authentication values from weakest to strongest:

1. none
Not secure; this is the default setting.
2. anonymous
This option is not secure against MITM attacks.
In-flight data is encrypted.
3. one-way
This method requires the use of certificates.
This is not secure against MITM attacks.
In-flight data is encrypted.
4. two-way password
This option is secure against MITM attacks.
In-flight data is encrypted.

5. two-way

This option requires the user of certificates.

This is the most secure option, and is secure against MITM attacks.

In-flight data is encrypted.

Note that "anonymous" and "one-way" are only allowed for per-client settings, not global settings.

Encryption settings

The following list ranks encryption values from weakest to strongest:

1. none

Not secure; this is the default setting.

Can only be specified if authentication is "none."

2. medium

Employs AES 128 and SHA-1.

3. high

Employs AES 256 and SHA-1.

Global authentication

Global authentication and encryption values can only be set through command-line interface (CLI) commands on the DD Boost Server. The CLI commands you use to set these values are described in the following sections.

For a complete list of DD Boost commands and options, see the *Data Domain Command Reference Guide for Version 6.1*

None

```
ddboost option set global-authentication-mode none
global-encryption-strength none
```

"None" 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 MITM attacks. Your system can operate in the same manner as before without suffering any performance degradation due to TLS.

When authentication is set to "none," encryption must be set to "none." If you select a different setting for authentication than "none," the encryption setting cannot be "none."

Two-way password

```
ddboost option set global-authentication-mode two-way-password
global-encryption-strength {medium | high}
```

The `two-way password` method performs two-way authentication using TLS with pre-shared key (PSK) authentication. Both the client and the Data Domain system are authenticated using the previously established passwords. When this option is selected, all data and messages between the client and the Data Domain system are encrypted.

This option is the only secure option available with DD Boost for OpenStorage and protects fully against man-in-the-middle (MITM) attacks.

Encryption strength must be either `medium` or `high`.

Two-way password authentication is unique because it is the only method that is both secure against MITM and can be done without the caller specifying it.

Two-way

```
ddboost option set global-authentication-mode two-way
global-encryption-strength {medium | high}
```

This is the most secure option.

The two-way option employs TLS with certificates. Two-way authentication is achieved using certificates provided by the application.

This setting is compatible with existing use of certificates. Setting the global authentication setting to "two-way" requires all applications that connect to the Data Domain system to support and supply certificates.

Any application that does not support certificates and does not specify two-way authentication and provide certificates through the `ddp_connect_with_config` API will fail.

Backwards Compatibility Scenarios

Older client and new Data Domain system

In this case, an application using a Boost library is employed with DD OS 6.1 or later. In this scenario, the client cannot perform two-way-password authentication, which has the following ramifications:

- Any global authentication settings must be set to "none" or "two-way" since the client cannot perform "two-way-password" authentication. Per-client authentication settings can be any value except "two-way-password" for the same reason.
- Any global or per-client settings of two-way password will cause applications with older client libraries to fail.
- The new Data Domain system will support existing connection protocols for old clients.

New client and older Data Domain system

The older Data Domain system cannot perform "two-way-password" authentication, which has the following ramifications:

- There are no global authentication or encryption settings.
- The per-client Data Domain system authentication setting cannot be "two-way password."
- The client will first attempt to use the new connection protocol or RPC; upon failure, the client reverts to the old protocol.
- The client can connect with other authentication methods except "two-way-password."

Authentication and Encryption Setting Examples

The following tables show examples in which settings are specified using calls, per-client settings, and global settings, and whether those settings can succeed.

These examples assume you have a DD Boost client connection to a Data Domain system with DD OS 6.1 or later. These examples do not apply to either of the situations described in [Backwards Compatibility Scenarios](#).

Note

If the global or per-client setting requires two-way authentication, the caller must specify it and provide the necessary certificates.

Table 3 One Setting

Call specifics	Per-client settings	Global settings	Used values
None	None	None	SUCCEEDS Authentication: none Encryption: none
Authentication: two-way-password Encryption: medium	None	None	SUCCEEDS Authentication: two-way-password Encryption: medium
None	Authentication: two-way-password Encryption: medium	None	SUCCEEDS Authentication: two-way-password Encryption: medium
None	None	Authentication: two-way-password Encryption: medium	SUCCEEDS Authentication: two-way-password Encryption: medium
None	None	Authentication: two-way Encryption: high	FAILS Two-way and high are required. The client must specify two-way and provide certificates.
Authentication: two-way Encryption: high	None	None	SUCCEEDS Authentication: two-way Encryption: high

Table 4 Multiple Settings

Call specifics	Per-client settings	Global settings	Used values
Authentication: two-way Encryption: medium	None	Authentication: two-way Encryption: high	FAILS Two-way and high are required.
None	Authentication: two-way Encryption: high	Authentication: two-way-password Encryption: medium	FAILS Two-way and high are required. The client must specify two-way and provide certificates.
Authentication: two-way Encryption: high	Authentication: two-way Encryption: high	Authentication: two-way Encryption: medium	SUCCEEDS Authentication: two-way Encryption: high

Table 4 Multiple Settings (continued)

Call specifies	Per-client settings	Global settings	Used values
None	Authentication: two-way-password Encryption: medium	Authentication: two-way Encryption: medium	FAILS Two-way and medium are required. The client must specify two-way and provide certificates.
Authentication: two-way Encryption: high	Authentication: two-way Encryption: medium	Authentication: two-way-password Encryption: medium	SUCCEEDS Authentication: two-way Encryption: high

Managed File Replication (MFR)

The DD Boost software enables applications to control the Data Domain Replicator software so that copies of data on one Data Domain system can be created on a second Data Domain system using the network-efficient Data Domain replication technology.

Because backup applications control replication of data between multiple Data Domain systems, they can provide backup administrators with a single point of management for tracking all backups and duplicate copies.

Dynamic interface groups provide the ability to control the interfaces used for DD Boost MFR, to direct the replication connection over a specific network, and to use multiple network interfaces with high bandwidth and reliability for failover conditions. For more information, see [Using Dynamic Interface Groups for MFR](#) on page 50.

Low-Bandwidth Optimization

The low-bandwidth Replicator option reduces the WAN bandwidth utilization. It is useful if managed file replication is being performed over a low-bandwidth network (WAN) link. This feature provides additional compression during data transfer and is recommended only for managed file replication jobs that occur over WAN links that have fewer than 6Mb/s of available bandwidth.

Both the source and destination Data Domain systems must be configured with this setting to enable low-bandwidth optimization, and the option applies to all replication jobs.

For more information about this topic, refer to the *EMC Data Domain Operating System Administration Guide*.

Encrypted Managed File Replication

This option allows applications to use SSL to encrypt the replication session between two Data Domain systems. All data and metadata is sent encrypted over the network.

The source and destination systems negotiate automatically to perform encryption transparent to the requesting application. Encrypted file replication uses the ADH-AES256-SHA cipher suite.

The option is enabled on each Data Domain system and applies to all managed file replication jobs on that system. Both the source and the destination Data Domain systems participating in managed file replication jobs must have this option enabled.

Encrypted managed file replication can be used with the encryption of data-at-rest feature available on the DD OS with the optional Encryption license. When encrypted managed file replication is used with the encryption of data-at-rest feature, the encrypted backup image data is encrypted again using SSL for sending over WAN.

Note

- For more information about this topic, see the *EMC Data Domain Operating System Administration Guide*. Both the source and the destination Data Domain systems must be running DD OS 5.0 or later to use this feature. Enabling this feature does not require restarting the file system on a Data Domain system.
 - The low-bandwidth optimization option and the encryption option can be used together.
-

DD Boost and High Availability

Beginning with DD OS 5.7.1, Data Domain systems with DD Boost can accommodate high availability (HA) configurations.

During normal operations, DD Boost on the Active node sends to the Standby node any Boost data and state information necessary to continue Boost operations on the Standby node if a failure should occur.

Note

DD Boost currently supports only Active-Standby configurations.

DD Boost performs periodic operations to force user data to disk on the server. Boost on the client buffers all user data between these periodic synchronize-to-disk operations so that if a DD server fails, the data can be re-sent.

This method applies to virtual writes as well. You can mix standard write operations with synthetic write operations.

DD Boost, HA, and failover

When a Data Domain system with HA enabled fails, recovery occurs in less than ten minutes. Once the failed system recovers, DD Boost recovery begins and applications using Boost automatically recover without failing or receiving an error. DD Boost recovery may take longer than ten minutes since Boost recovery cannot begin until failover of the DD system is complete.

No changes are necessary to allow applications to take advantage of DD Boost HA capabilities. When using DD Boost 3.2.1 and DD OS 5.7.1 on HA configurations, applications automatically recover if a failover occurs. No action is required from the application.

Partial HA configurations

Managed File Replication (MFR) is supported between any two Data Domain systems running compatible versions of DD OS, regardless of whether one or both of the DD systems is enabled for HA.

MFR between two HA systems will succeed in the event of failure of either system since both support HA. An MFR in progress will recover seamlessly if either the source HA system or the destination HA system fails.

In addition, MFR between an HA system and a non-HA system will succeed if the HA system fails; it will not succeed if the non-HA system fails.

MFR to HA-enabled systems

A single node DD system running DD OS 5.7 or later and performing MFR to an HA system recovers seamlessly if the HA system fails. The MFR will not recover seamlessly if the single node DD source system fails.

MFR from HA-enabled systems

An MFR from an HA system to a single-node Data Domain system running DD OS 5.7 or later recovers seamlessly if the source HA system fails. However, the MFR will not recover seamlessly if the single- node DD destination system fails.

Note that in all cases involving partial HA configurations, the non-HA system must be running DD OS 5.7 to allow an MFR to continue seamlessly should a failure occur. In partial HA configurations where the non-HA system is running a version of DD OS older than 5.7, the MFR will not recovery seamlessly from a failure of either system.

In all cases, the application must be using DD HA Boost 3.2.1 libraries for seamless recovery of MFR to occur.

MTree Replication

Beginning with DD OS Release 5.5, MTree replication for storage units is supported for different user names on source and destination Data Domain systems. To enable MTree replication for storage units, you must convert the target storage unit from MTree to storage unit by assigning a user to the MTree. To assign a user to the MTree, use the DD OS `ddboost storage-unit modify` command. (See the *EMC Data Domain Operating System Command Reference Guide* for details.)

IPv6 Support

IPv6 replication support includes managed file replication, which you configure using the `ddboost file-replication option set ipversion ipv6` command.

The client connects to the Data Domain system using the hostname. The hostname parameter is of type string and can also accept an IPv4 address in the form a.b.c.d or any valid IPv6 address (1234:abcd::4567 or 12:34:56:78::0, for example). If both IPv4 and IPv6 addressing exist in the network, the IP address family that is provided by the client upon connection is used as the preferred IP address family to resolve the hostname. If a single IP address family exists in the network (only IPv4 or only IPv6), then the hostname resolves to that address, and that address is used for the client-to-Data Domain backup and restore connection. If no preferred IP address family is specified by the client, then the client-to-Data Domain backup and restore connection will use whatever IP address that the DNS resolves. The default is IPv4. For backward

compatibility, IPv4 is set as the preferred IP address. If the address resolution fails, it is up to the client to try to reconnect with a new hostname.

Dynamic Interface Groups: DD Boost IP Data Path Management

Note

This feature applies to the DD Boost-over-IP transport only.

The Dynamic Interface Groups (DIG) feature lets you combine multiple Ethernet links into a group and register only one interface on the Data Domain system with the backup application. The DD Boost Library negotiates with the Data Domain system on the interface registered with the application to obtain the best interface to send data to the Data Domain system. Load balancing provides higher physical throughput to the Data Domain system compared to configuring the interfaces into a virtual interface using Ethernet-level aggregation (using LACP, for example).

The Data Domain system load balances the connections coming in from multiple backup application hosts on all interfaces in the group. Load balancing is transparent to the backup application and is handled by the DD Boost software. Because DIG works at the DD Boost software layer, it is seamless to the underlying network connectivity and supports physical and virtual interfaces. The data transfer is load-balanced based on the number of connections outstanding on the interfaces. Only connections for backup and restore jobs are load-balanced.

DIG also works with other network layer functionality on Data Domain systems, including VLAN tagging and IP aliasing. This functionality allows additional flexibility in segregating traffic into multiple virtual networks, all of which run on the same physical links on the Data Domain system.

Note

See the *EMC Data Domain Operating System Administration Guide* for more information about how to configure VLAN tagging and IP aliasing on a Data Domain system.

DIG also provides the ability to control the interfaces used for DD Boost MFR, to direct the replication connection over a specific network, and to use multiple network interfaces with high bandwidth and reliability for failover conditions. For more information, see [Using Dynamic Interface Groups for MFR](#) on page 50.

The DIG feature provides the following benefits:

- Eliminates the need to register the Data Domain system on multiple interfaces with the application, which simplifies installation and configuration.
- Transparently fails over all in-process jobs from the failed interface to healthy operational links. From the point of view of the backup application, the jobs continue uninterrupted.
- Routes subsequent incoming backup jobs to the available interfaces if one of the interfaces in the group goes down while the Data Domain system is still operational.
- Automatically load-balances backup and restore jobs on multiple interfaces in the group, resulting in higher utilization of the links.
- Works with 1-GbE interfaces and 10-GbE interfaces in the same interface group. Combining interfaces of different speeds in a single DIG is allowed and supported.

- An administrator can define multiple interface groups where load balancing and failover apply within a DIG *<group-name>*. This increases the capability to support a backup server that can reach only some of the Data Domain system interfaces, such as clients on VLANs.
- Each interface group *<group-name>* includes a list of interfaces and clients that belong to the DIG. Within a DIG *<group-name>*, all interfaces are reachable by all the clients for *<group-name>*.
- Public IP-to-private VLAN configuration using client host range:
 - Clients can reach the Data Domain private network if they are on the same subnet.
 - Avoids static route on clients by adding IP alias/VLAN IP on the Data Domain system to match the client subnet.
 - Clients on the same domain name need to reach different private networks—the alias/VLAN IP network.
 - Redirects clients off the public network to the appropriate private network for data isolation or to avoid configuration of static routes, keeping the client and Data Domain IP addresses on the same subnet.

For more information, see [Clients](#) on page 27.

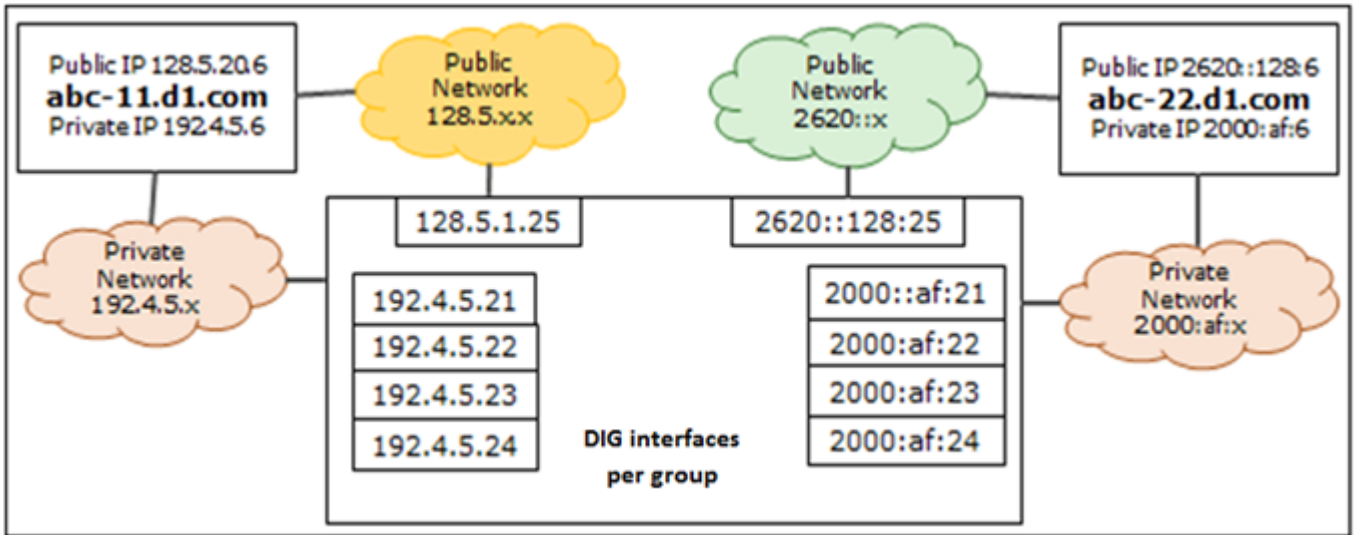
Interfaces

A Dynamic Interface Group (DIG) interface is a member of a single interface group *<group-name>* and may have the following characteristics:

- Physical interface such as `eth0a`
- Virtual interface, created for link failover or link aggregation, such as `veth1`
- Virtual alias interface such as `eth0a:2` or `veth1:2`
- Virtual VLAN interface such as `eth0a.1` or `veth1.1`
- Within an interface group *<group-name>*, all interfaces must be on unique interfaces (Ethernet, virtual Ethernet) to ensure failover in the event of network error.

DIG provides full support for static IPv6 addresses, providing the same capabilities for IPv6 as for IPv4. Concurrent IPv4 and IPv6 client connections are allowed. A client connected with IPv6 sees IPv6 DIG interfaces only. A client connected with IPv4 sees IPv4 DIG interfaces only. Individual interface groups include all IPv4 addresses or all IPv6 addresses.

Figure 3 DIG Support for IPv4 and IPv6 Addressing



Interface Enforcement

The Dynamic Interface Group (DIG) feature gives you the ability to 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 DIG 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
```

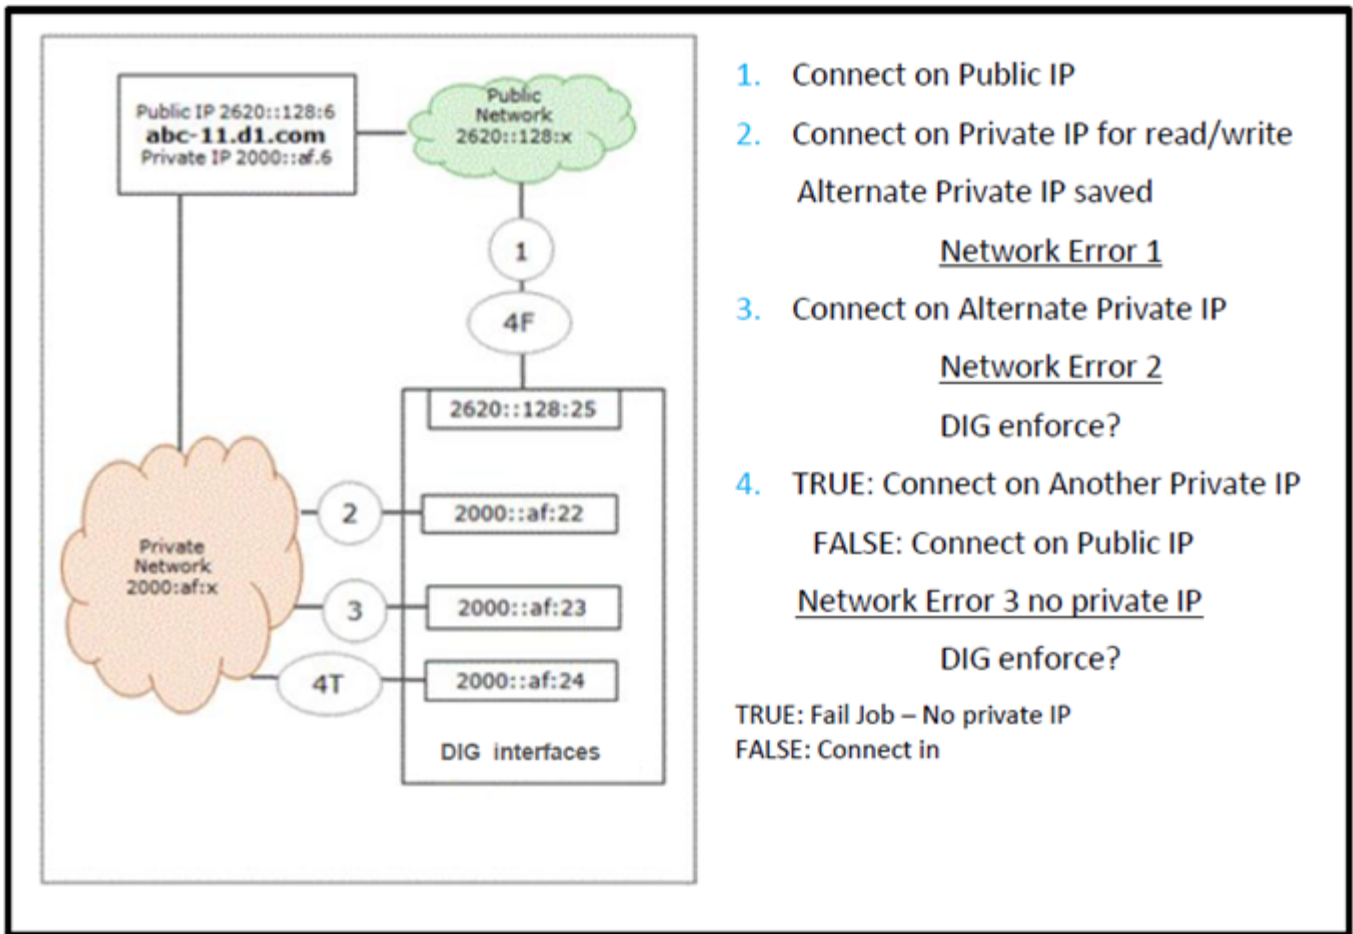
Note

In the previous example, `ifgroup` refers to the interface group.

After you've made this entry in the registry, you must do a `filesys restart` for the setting to take effect.

The following illustration shows the decision flow for DIG connections. If interface enforcement is on (TRUE), the system always attempts to reconnect on a private IP address when a job fails. If a private IP address is not available, the job is canceled, and a `Cancel job for non-ifgroup interface error` message is generated. If interface enforcement is off (FALSE), a failed job resumes using a public IP address.

Figure 4 DIG Connection Decision



Clients

A Dynamic Interface Group (DIG) client is a member of a single interface group `<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 for the client, such as `ddboost`
- Client public IP range, such as `128.5.20.0/24`

Prior to write or read processing, the client requests a DIG IP address from the server. To select the client DIG association, the client information is evaluated according to the following order of precedence (see [Figure 5](#) on page 28):

1. IP address of the connected Data Domain system. If there is already an active connection between the client and the Data Domain system, and the connection exists on the interface in the DIG, then the DIG interfaces are made available for the client.
2. Connected client IP range. An IP mask check is done against the client source IP; if the client's source IP address matches the mask in the DIG clients list, then the DIG interfaces are made available for the client.
 - For IPv4, `xx.xx.xx.0/24` (`128.5.20.0/24` in [Figure 5](#) on page 28) provides a 24-bit mask against the connecting IP. The `/24` represents which

bits are masked when the client's source IP address is evaluated for access to the DIG. For IPv4, 16, 20, 24, 28, and 32 bit masks are supported.

- For IPv6, `xxxx::0/112` provides a 112-bit mask against the connecting IP. The /112 represents what bits are masked when the client's source IP address is evaluated for access to the DIG. For IPv6, 64, 112, and 128 bit masks are supported.
- DD OS 6.1 supports any five masks ranging from /8 to /32.
- DD OS 6.1 supports prefix hostname matches (for example, `abc_.emc.com` for hosts that have the prefix `abc`).

This host-range check is useful for separate VLANs with many clients where there isn't a unique partial hostname (domain).

3. Client Name: `abc-11.d1.com`

4. Client Domain Name: `*.d1.com`

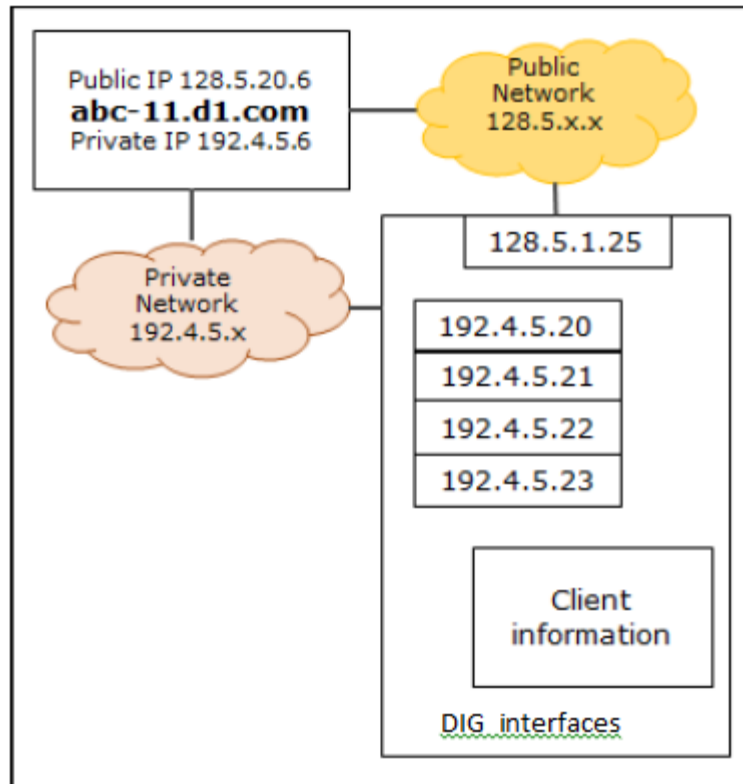
5. All Clients: `*`

Note

In a mixed network with IPv4 and IPv6 addressing, DIG client configuration should not allow IPv6 to match an IPv4 group, nor should it allow IPv4 to match an IPv6 group. Therefore, "*" should not be configured. Also, if the clients on IPv4 and IPv6 are on the same domain name (`*.domain.com`, for example), only fully qualified domain names or host-range (IP with mask) should be used.

If none of these checks find a match, DIG interfaces are not used for this client.

Figure 5 DIG Host Range for Client Selection



Using interface groups for Managed File Replication (MFR)

Interface groups can be used to control the interfaces used for DD Boost MFR, to direct the replication connection over a specific network, and to use multiple network interfaces with high bandwidth and reliability for failover conditions. All Data Domain IP types are supported—IPv4 or IPv6, Alias IP/VLAN IP, and LACP/failover aggregation.

Note

Interface groups used for replication are different from the interface groups previously explained and are supported for DD Boost Managed File Replication (MFR) only. For detailed information about using interface groups for MFR, see the *EMC Data Domain Boost for Partner Integration Administration Guide* or the *EMC Data Domain Boost for OpenStorage Administration Guide*.

Without the use of interface groups, configuration for replication requires several steps:

1. Adding an entry in the `/etc/hosts` file on the source Data Domain system for the target Data Domain system and hard coding one of the private LAN network interfaces as the destination IP address.
2. Adding a route on the source Data Domain system to the target Data Domain system specifying a physical or virtual port on the source Data Domain system to the remote destination IP address.
3. Configuring LACP through the network on all switches between the Data Domain systems for load balancing and failover.
4. Requiring different applications to use different names for the target Data Domain system to avoid naming conflicts in the `/etc/hosts` file.

Using interface groups for replication simplifies this configuration through the use of the DD OS System Manager or DD OS CLI commands. Using interface groups to configure the replication path lets you:

- Redirect a hostname-resolved IP address away from the public network, using another private Data Domain system IP address.
- Identify an interface group based on configured selection criteria, providing a single interface group where all the interfaces are reachable from the target Data Domain system.
- Select a private network interface from a list of interfaces belonging to a group, ensuring that the interface is healthy.
- Provide load balancing across multiple Data Domain interfaces within the same private network.
- Provide a failover interface for recovery for the interfaces of the interface group.
- Provide host failover if configured on the source Data Domain system.
- Use Network Address Translation (NAT)

The selection order for determining an interface group match for file replication is:

1. Local MTree (storage-unit) path and a specific remote Data Domain hostname
2. Local MTree (storage-unit) path with any remote Data Domain hostname
3. Any MTree (storage-unit) path with a specific Data Domain hostname

The same MTree can appear in multiple interface groups only if it has a different Data Domain hostname. The same Data Domain hostname can appear in multiple interface

groups only if it has a different MTree path. The remote hostname is expected to be an FQDN, such as dd890-1.emc.com.

The interface group selection is performed locally on both the source Data Domain system and the target Data Domain system, independent of each other. For a WAN replication network, only the remote interface group needs to be configured since the source IP address corresponds to the gateway for the remote IP address.

IP Failover Hostname

The Failover Hostname feature lets you configure an alternative Data Domain administrative IP address and hostname for use on failover at first connection or on failover resulting from network errors. You can configure the alternative hostname in DNS or in the `/etc/hosts` file on the DD Boost client. Both IPv4 and IPv6 are supported.

To configure the alternative hostname, append `-failover` to the Data Domain system hostname.

IPv4 Example:

```
10.6.109.38 ddp-880-1.datadomain.com ddp-880-1
10.6.109.40 ddp-880-1-failover.datadomain.com ddp-880-1-failover
```

IPv6 Example:

```
3000::230 ddp-880-2-v6.datadomain.com ddp-880-2-v6
3000::231 ddp-880-2-v6-failover.datadomain.com ddp-880-2-v6-failover
```

This feature eliminates the need to have the administrative IP address in link failover mode. In addition, you can add this failover interface to an interface group so you can connect directly to this group without going through the system's standard administrative interface, thereby improving load balance and throughput performance. If the initial connection fails, the failover IP address is used, if it is available. Once the connection is established, interface group is used to select the read/write interfaces. Using the IPv4 example above:

1. The client attempts to connect to `ddp-880-1.datadomain.com`.
2. If the connection fails, the client attempts to connect to `ddp-880-1-failover.datadomain.com`.
3. If network errors occur after the initial connection is made, the connection is retried on the other interface. If the initial connection was on `ddp-880-1-failover.datadomain.com`, for example, the client retries the connection on `ddp-880-1.datadomain.com`. The last address attempted on errors is always the Data Domain system IP address.


Note

On Windows 2008 R2, the "TcpTimedWaitDelay" registry entry for timing out connections may be missing. This registry entry is essential to allow host-failover recovery. The name of the registry key in windows 2008 R2 is: `HKLM\System\CurrentControlSet\Services\Tcpip\Parameters`. This key should be set to a value of: double word "10".

DD Boost-over-Fibre Channel Transport

In earlier versions of DD OS, all communication between the DD Boost Library and any Data Domain system was performed using IP networking. The application specified the Data Domain system using its hostname or IP address. See [Figure 6](#) on page 31.

Figure 6 DD Boost-over-IP Transport



1. Backup Server
2. Applications, DD Boost Library, TCP/IP Transport
3. Data Domain System
4. DD Boost Service
5. TCP/IP

DD OS now offers an alternative transport mechanism for communication between the DD Boost Library and the Data Domain system — Fibre Channel.


Note

Windows, Linux, HP-UX, AIX, and Solaris client environments are supported.

To request access to a Data Domain system using the DD Boost-over-FC transport, the application specifies the Data Domain system using the special string `DFC-<dfc-server-name>`, where `<dfc-server-name>` is the DD Boost-over-FC server name configured for the Data Domain system.

Note

Just as IP hostnames are not case-sensitive, the `dfc-server-name` is not case-sensitive.

Figure 7 SCSI Commands between Backup Server and Data Domain system.


1. Backup Server
2. Application, DD Boost Library, DD Boost-over-FC Transport
3. Data Domain System
4. DD Boost Service
5. DD Boost-over-FC Server
6. SCSI Commands over FC
7. SCSI Processor Devices

Setting up the DD Boost-over-FC service on the Data Domain system requires additional configuration steps. See [Configuring DD Boost-over-FC Service](#) for details.

For the DD Boost-over-FC transport, load balancing and link-level high availability is achieved through a different means, not through Dynamic Interface Groups (DIG). See the section [DD Boost-over-Fibre Channel Path Management](#) for a description.

Note

The DD Boost-over-FC communication path applies only between the backup server/DD Boost Library and the Data Domain system, and does not apply to communication between two Data Domain systems. As shown in the next figure, such communication is ALWAYS over an IP network, regardless of the communication path between the backup server and the Data Domain systems.

Figure 8 Fibre Channel Communication Path

1. Backup Server
2. Application, DD Boost Library
3. IP or FC
4. IP or FC (Control)
5. Data Domain System, Replication Source
6. IP ONLY (Data)
7. Data Domain System, Replication Destination

DD Boost-over-Fibre Channel Path Management

The Dynamic Interface Group (DIG)-based mechanism described in [DIG: DD Boost IP Load Balancing and Failover](#) is based on Ethernet interfaces and is not applicable to the Fibre Channel transport. Instead, a different path mechanism is provided for the DD Boost-over-FC solution.


The Data Domain system advertises one or more Processor-type SCSI devices to the backup server, over one or more physical paths. The operating system discovers all devices through all available paths, and creates a generic SCSI device for each discovered device and path.

For example, consider the case where:

- Backup server has 2 initiator HBA ports (A and B)
- Data Domain System has 2 FC target endpoints (C and D)
- Fibre Channel Fabric zoning is configured such that both initiator HBA ports can access both FC target endpoints

- Data Domain system is configured with a SCSI target access group containing:
 - Both FC target endpoints on the Data Domain System
 - Both initiator HBA ports
 - 4 devices (0, 1, 2, and 3)

Figure 9 DD Boost-over-FC Path Management Scenario



1. Four Devices
2. Backup Server
3. HBA Initiator A
4. HBA Initiator B
5. Data Domain System
6. Fibre Channel Endpoint C
7. Fibre Channel Endpoint D

In this case, the backup server operating system may discover up to 16 generic SCSI devices, one for each combination of initiator, target endpoint, and device number:

- /dev/sg11: (A, C, 0)
- /dev/sg12: (A, C, 1)
- /dev/sg13: (A, C, 2)
- /dev/sg14: (A, C, 3)
- /dev/sg15: (A, D, 0)
- /dev/sg16: (A, D, 1)
- /dev/sg17: (A, D, 2)
- /dev/sg18: (A, D, 3)
- /dev/sg19: (B, C, 0)
- /dev/sg20: (B, C, 1)
- /dev/sg21: (B, C, 2)
- /dev/sg22: (B, C, 3)
- /dev/sg23: (B, D, 0)
- /dev/sg24: (B, D, 1)
- /dev/sg25: (B, D, 2)
- /dev/sg26: (B, D, 3)

When the application requests that the DD Boost Library establish a connection to the server, the DD Boost-over-FC Transport logic within the DD Boost Library uses SCSI requests to build a catalog of these 16 generic SCSI devices, which are paths to access the DD Boost-over-FC service on the desired Data Domain System. As part of

establishing the connection to the server, the DD Boost-over-FC Transport logic provides to the server this catalog of paths.

Initial Path Selection

The server maintains statistics on the DD Boost-over-FC traffic over the various target endpoints and known initiators. During the connection setup procedure, Path Management logic in the server consults these statistics, and selects the path to be used for this connection, based upon the following criteria:

- For Queue-Depth Constrained clients (see below), evenly distribute the connections across different paths
- Choose the least busy target endpoint
- Choose the least busy initiator from among paths to the selected target endpoint

Dynamic Re-Balancing

The server periodically performs dynamic re-balancing. This involves consulting the statistics to look for situations where:

- For Queue-Depth Constrained clients (see below), connections are distributed unequally across available paths
- Workload across target endpoints is out of balance
- Workload across initiators is out of balance

If such a situation is discovered, the server may mark one or more connections for server-directed path migration. This is achieved by having the server request, during a future data transfer operation, that the DD Boost Library start using a different available path from the catalog for subsequent operations.

Client Path Failover

The client may start using a different path because it is directed to do so by the server dynamic re-balancing logic. But the client may also decide, on its own, to start using a different available path. This happens if the client receives errors when using the connection's current path.

For example, assume the path catalog for a connection consists of 8 paths:

- /dev/sg21: (A, C, 0)
- /dev/sg22: (A, C, 1)
- /dev/sg23: (A, D, 0)
- /dev/sg24: (A, D, 1)
- /dev/sg25: (B, C, 0)
- /dev/sg26: (B, C, 1)
- /dev/sg27: (B, D, 0)
- /dev/sg28: (B, D, 1)

and the server selects the (A, C, 0) path during initial path selection. The DFC transport logic in the DD Boost Library starts sending and receiving data for the connection, using SCSI commands to /dev/sg21.

Later, the link from target endpoint C to its switch becomes unavailable, due to cable pull or some hardware failure. Any subsequent SCSI request submitted by the DFC transport logic to /dev/sg21 will fail with an error code indicating that the SCSI request could not be delivered to the device.

In this case, the DFC transport logic looks in the catalog of devices, for a path with a different physical component; that is, a different combination of initiator and target endpoint. The SCSI request is retried on the selected path, and the process is repeated until a path is discovered over which the SCSI request can be successfully completed.

Queue-Depth Constraints

For the purposes of the DD Boost-over-FC solution, the specific SCSI device over which a request is received is irrelevant. All SCSI devices are identical, destination objects for SCSI commands as required by the SCSI protocol. When processing a SCSI request, the server logic gives no consideration to the specific device on which the SCSI request arrived.

Why bother to allow for more than one device? Because certain client-side operating systems impose a restriction on the number of outstanding IO requests which can be conducted simultaneously over a given generic SCSI device. For example, the Windows SCSI Pass-Through Interface mechanism will only conduct 1 SCSI request at a time through each of its generic SCSI devices. This impacts the performance of the DD Boost-over FC solution, if multiple connections (e.g. backup jobs) are trying to use the same generic SCSI device.

Additionally, the Data Domain system also imposes a limit on the number of outstanding IO requests per advertised SCSI device. For performance reasons with larger workloads, multiple SCSI devices may need to be advertised on the Data Domain system.

We use the term “queue-depth” to describe the system-imposed limit on the number of simultaneous SCSI requests on a single device. Client systems (like Windows) whose queue depth is so low as to impact performance are considered “queue-depth constrained.”

Refer to [Sizing DD Boost-over FC Device-Set](#) on page 64 for guidance regarding how many devices to configure based on the workload, type of Data Domain system, and whether or not the client system is queue-depth constrained.

Virtual Synthetic Backups

A synthetic full or synthetic cumulative incremental backup is a backup assembled from previous backups. Synthetic backups are generated from one previous, traditional full or synthetic full backup, and subsequent differential backups or a cumulative incremental backup. (A traditional full backup means a non-synthesized, full backup.) A client can use the synthesized backup to restore files and directories in the same way that a client restores from a traditional backup.

During a traditional full backup, all files are copied from the client to a backup server and the resulting image set is sent to the Data Domain system. The files are copied even though those files may not have changed since the last incremental or differential backup. During a synthetic full backup, the previous full backup and the subsequent incremental backups on the Data Domain system are combined to form a new, full backup. The new, full synthetic backup is an accurate representation of the clients' file system at the time of the most recent full backup.

Because processing takes place on the Data Domain system under the direction of the backup server instead of the client, virtual synthetic backups help to reduce the network traffic and client processing. Client files and backup image sets are transferred over the network only once. After the backup images are combined into a synthetic backup, the previous incremental and/or differential images can be expired.

The virtual synthetic full backup is a scalable solution for backing up remote offices with manageable data volumes and low levels of daily change. If the clients experience a high rate of daily change, the incremental or differential backups are too large. In this case, a virtual synthetic backup is no more helpful than a traditional full backup. To ensure good restore performance it is recommended that a traditional full backup be created every two months, presuming a normal weekly full and daily incremental backup policy.

The virtual synthetic full backup is the combination of the last full (synthetic or full) backup and all subsequent incremental backups. It is time stamped as occurring one second after the latest incremental backup. It does NOT include any changes to the backup selection since the latest incremental backup.

Client Access Validation

Configuring client access validation for DD Boost limits access to the Data Domain system for DD Boost clients by requiring DD Boost authentication (per connection) for:

- the initial connection to the Data Domain system
- each restart of DD Boost (Enable/Disable)
- each file system restart
- each Data Domain system reboot

The list of clients can be updated at anytime without a restart requirement, thus eliminating access validation impact on jobs in progress.

DD Boost Multiuser Data Path

DD Boost multiuser data path enhancements improve storage unit isolation. Multiple users can be configured for DD Boost access on a Data Domain system.

Storage Unit Management

You can use DD OS `ddbboost` commands to configure and modify storage units, tenants, and quota limits, and to configure stream warning limits for each storage unit.

Multiuser Storage Units Access Control

The Multiuser Storage Unit Access Control feature for DD Boost 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 cryptography.

The system administrator creates a local Data Domain user for each backup application to be used for their storage units. The storage unit user is required when the storage unit is created. When backup applications connect to the Data Domain system, the applications can only access the storage units owned by the username used to make the connection. Access to a storage unit is determined dynamically so that changes to a storage unit's username take effect immediately. When a storage unit's username is changed to another username, all read and write operations by the backup application using the old username fail immediately with permission errors.

The `tenant-unit` keyword is introduced to the `ddbboost storage-unit` command for integration with the Data Domain Secure Multi-Tenancy feature. One storage unit must be configured for each tenant unit. Each tenant unit can be associated with multiple storage units. Tenant unit association and storage unit username ownership are independent from each other. The tenant unit is used for management path using the command-line-interface, but cannot be used for data path, for example, read and write. All commands for storage units support tenant units.

Note

For more information about tenant units, refer to the *EMC Data Domain Operating System Administration Guide*.

Storage Unit Capacity Quotas

DD OS users can use quotas to provision Data Domain system logical storage limits, ensuring that dedicated portions of the Data Domain system are available as unique storage units. DD Boost storage-unit quota limits may be set or removed dynamically. Quotas may also be used to provision various DD Boost storage units with different logical sizes, enabling an administrative user to monitor the usage of a particular storage unit over time.

You can also configure the reported physical size; this is the size reported to the backup application. The physical size is the Disk Pool "raw size" in NetBackup. On the Data Domain system itself, the actual size is shown. The logical capacity quota is still available if you configure the physical size. You can modify the reported physical size at a later time using `ddbboost storage-unit modify`. You can display the reported physical size using `ddbboost storage-unit show`.

See the `ddbboost`, `quota`, and `mtree` sections of the *EMC Data Domain Operating System Command Reference Guide* for details on the quota feature, and commands pertaining to quota operations.

Note

Be careful with this feature when you are using backup applications that use the DD Boost API for capacity management. The DD Boost API attempts to convert the logical setting to a physical setting for the API by dividing the logical setting by the deduplication ratio. Logical quotas may need to be adjusted when the deduplication ratio changes.

Storage Units Stream Count Management

You can configure five types of stream warning limits for each storage unit:

- `write-stream-soft-limit`
- `read-stream-soft-limit`
- `repl-stream-soft-limit`
- `combined-stream-soft-limit`
- `combined-stream-hard-limit`

For each storage unit, stream counters are maintained to monitor backup, restore, replication-in, and replication-out data. To configure stream limits when creating a storage unit, use the `ddbboost storage-unit create` command. To configure stream limits for an existing storage unit, use the `ddbboost storage-unit modify`

command. To display the active streams per storage unit, use the `ddboost streams show active` command.

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.

Any of these stream warning limits can also be set to `none`.

Note

DD Boost backup applications are expected to reduce their workload to remain below the stream warning quotas. You can reconfigure the warning limit to avoid exceeding the quotas.

For more information about configuring stream limits, see [Configuring Storage Units with Stream Limits \(Optional\)](#) on page 45.

CHAPTER 3

Preparing the Data Domain System for DD Boost

Note

Complete descriptions of commands used in this guide are provided in the *EMC Data Domain Operating System Command Reference Guide*.

This chapter covers the following topics:

- [Enabling DD Boost on a Data Domain System](#)..... 42
- [Assigning Multiple Users to DD Boost](#)..... 42
- [Creating Storage Units](#) 43
- [Configuring Logical Quotas for Storage Units \(Optional\)](#) 44
- [Configuring Storage Units with Stream Limits \(Optional\)](#)..... 45
- [Configuring Distributed Segment Processing](#)..... 46
- [Configuring Dynamic Interface Groups](#) 47
- [Using Dynamic Interface Groups for MFR](#)..... 50
- [Configuring MFR](#)..... 60
- [Configuring Client Access Validation](#)..... 61
- [Configuring DD Boost-over-FC Service](#)..... 62
- [Setting Global Authentication and Encryption](#)..... 70

Enabling DD Boost on a Data Domain System

Every Data Domain system that is enabled for Data Domain Boost deduplication must have a unique name. You can use the DNS name of the Data Domain system, which is always unique.

Procedure

1. On the Data Domain system, log in as an administrative user.
2. Verify that the file system is enabled and running by entering:

```
# fileysys status
The file system is enabled and running.
```

3. Add the DD Boost license using the license key that Data Domain provided:

```
# license add license_code
License "ABCE-BCDA-CDAB-DABC" added.
```

4. Enable DD Boost deduplication by entering:

```
# ddbboost enable
DD Boost enabled
```

Note

- The users must be configured in the backup application to connect to the Data Domain system. For more information, refer to the *EMC Data Domain Operating System Administration Guide*.
- Multiple users can be configured for DD Boost access on a Data Domain system. The username, password, and role must have already been set up on the Data Domain system using the DD OS command:

```
user add <user> [password <password>]
[role {admin | security | user | backup-operator | data-access}]
[min-days-between-change <days>] [max-days-between-change
<days>]
[warn-days-before-expire <days>] [disable-days-after-expire
<days>]
[disable-date <date>]
```

For example, to add a user with a login name of `jsmith` and a password of `usr256` with administrative privilege, enter:

```
# user add jsmith password usr256 role admin
```

Then, to add `jsmith` to the DD Boost user list, enter:

```
# ddbboost user assign jsmith
```

Assigning Multiple Users to DD Boost

As system administrator, you need to create a local Data Domain system user for each backup application to use with their storage units. The storage units are either created with a username, or can be modified to change the username for an upgrade. Storage units are accessible only to applications with the username that owns the storage unit. Each storage unit is owned by one username, and the same username can own multiple storage units. The application passes the username and password to DD Boost, and DD Boost passes them to the Data Domain system when attempting to connect to the Data Domain system. The Data Domain system then authenticates the

username and password. The username and password can be shared by different applications.

When a storage unit is created with a valid Data Domain system local user but not assigned to DD Boost, the user is automatically added to the DD Boost users list in the same way that a user is added via the `ddboost user assign` command. If a storage unit is created without specifying the owning username, the current DD Boost user name is assigned as owner.

To assign and add one or more users to the DD Boost users list, enter:

```
# ddbboost user assign user1 user2
User "user1" assigned to DD Boost.
User "user2" assigned to DD Boost.
```

To verify and display the users in the users list, enter:

```
# ddbboost 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	-

To unassign and delete the user from the users list, enter:

```
# ddbboost user unassign user1
User "user1" unassigned from DD Boost.
```

Creating Storage Units

You need to create one or more storage units on each Data Domain system enabled for DD Boost.

Procedure

1. To create a storage unit on the Data Domain system, enter:

```
# ddbboost storage-unit create NEW_STU1 user user1
Created storage-unit "NEW_STU1" for "user1".
```

Note

A storage unit name must be unique on any given Data Domain system. However, the same storage unit name can be used on different Data Domain systems. The username owns the storage unit and ensures that only connections with this username's credentials are able to access this storage unit.

See the section on `ddboost storage-unit` in the *EMC Data Domain Operating System Command Reference Guide* for details on command options.

Note

When a storage-unit is created with a valid Data Domain local user who is not already assigned to DD Boost, the user is automatically added to the DD Boost user list in the same way that a `ddboost` user is added to the user list via the `ddboost user assign` command.

2. Repeat the above step for each Boost-enabled Data Domain system.
3. To modify a storage unit on the Data Domain system, enter:

```
# ddbboost storage-unit modify NEW_STU1 user user2
Storage-unit "NEW_STU1" modified for user "user2".
```

Note

The `ddboost storage-unit modify` command allows the backup application to change the user-name ownership of the storage unit. Changing the username does not need to change attributes of every file on the storage unit, therefore it is fast.

4. To display the users list for the storage units, enter:

```
# ddbboost 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       ddu1          -
getevent            0.0             RW      ddu1          -
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
```

Configuring Logical Quotas for Storage Units (Optional)

The storage on a Data Domain system can be provisioned through optional quota limits for a storage-unit. Quota limits can be specified either at the time of creation of a storage-unit, or later through separate commands. For more information, refer to the sections on quotas and `ddboost` in the *EMC Data Domain Operating System Command Reference Guide*.

Procedure

1. To enable quota limits on the Data Domain system, enter:


```
# quota enable
```
2. To configure quota limits at the time of creation of a storage unit, specify the `quota-soft-limit` and `quota-hard-limit` values with the following command:

```
# ddbboost storage-unit create storage-unit
[quota-soft-limit n {MiB|GiB|TiB|PiB}] [quota-hard-limit n {MiB|
GiB|TiB|PiB}]
```

3. To modify quota limits after they've been created, specify the new `quota-soft-limit` and `quota-hard-limit` values with the following command:

```
# ddbboost storage-unit modify storage-unit
[quota-soft-limit {n {MiB|GiB|TiB|PiB}|none}] [quota-hard-limit
{n {MiB|GiB|TiB|PiB}|none}]
```

4. To verify the quota limits of a storage unit:

```
# quota show storage-units storage-unit-list
```

Configuring Storage Units with Stream Limits (Optional)

The system administrator configures stream warning limits against each storage-unit for each of the five limits:

- `write-stream-soft-limit`
- `read-stream-soft-limit`
- `repl-stream-soft-limit`
- `combined-stream-soft-limit`
- `combined-stream-hard-limit`

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, you could enter:

```
# ddbboost 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 a storage unit, you could enter:

```
# ddbboost 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
```

Setting a limit to `none` disables that limit.

To display the DD Boost stream limits for all the active storage units, enter `ddbboost streams show active`. To display the DD Boost stream limits for a specific storage unit, enter:

```
# ddbboost 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 Data Domain system stream limits displayed in the output are based on the type of the Data Domain system.

To display the DD Boost stream limits history for a specific storage unit for a specific time, enter:

```
# ddbboost streams show history storage-unit NEW_STU0 last 1hours
INTERVAL: 10 mins
"- " indicates that the data is not available for the intervals
```

```
Storage-Unit: "NEW_STU0"
```

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

Configuring Distributed Segment Processing

The distributed segment processing option is configured on the Data Domain system and applies to all the backup servers and the DD Boost libraries installed on them.

The option can be configured using the following command:

```
# ddbboost option set distributed-segment-processing {enabled | disabled}
```

Note

Enabling or disabling the distributed segment processing option does not require a restart of the Data Domain file system.

Distributed segment processing is supported with version 2.2 or later of the DD Boost libraries communicating with a Data Domain system that is running DD OS 4.8 or later.

Distributed segment processing is enabled by default on a system initially installed with DD OS 5.2. If a system is upgraded from DD OS 5.1, 5.0.x or 4.9.x to DD OS 5.2, distributed segment processing is left in its previous state.

Distributed segment processing is enabled (and cannot be disabled) on DD OS 5.5.1.0 and earlier.

Distributed segment processing is enabled by default for Solaris plug-ins running on a SPARC T4 or T5 processor and running Solaris 11 (with SRU2 or later) or Solaris 11.1 or later.

Configuring Dynamic Interface Groups

Note

This feature applies only to DD Boost over IP. For an overview of the Dynamic Interface Group (DIG) feature, see [Dynamic Interface Groups: DD Boost IP Load Balancing and Failover](#).

When a Data Domain system receives a connection request from a client in a configured interface group, the DIG feature assigns the connection to the least used interface in the group, providing load balancing and higher input/output throughput.

To configure DIG, create an interface group on the Data Domain system by adding existing interfaces to the group as described below.

Procedure

1. Create the interface group:

```
# ifgroup create group_name
```

Examples:

```
# ifgroup create external
# ifgroup create lab10G
```

Note

The *group_name* “default” can be used without being created first. In all the remaining `ifgroup` commands, the “default” group is used if not specified.

2. Add clients and interfaces to each interface group. The interfaces must already have been created with the `net` command.

```
# ifgroup add group_name
{interface {ipaddr | ipv6addr} | client host}
```

This command provides full interface group 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 interface group interfaces only. A client connected with IPv4 sees IPv4 interface-group interfaces only. Individual interface groups include all IPv4 addresses or all IPv6 addresses.

Examples:

```
# ifgroup add interface 10.6.109.140 client *.datadomain.com
# ifgroup add interface 10.6.109.141 client *

# ifgroup add ipv4-group interface 192.4.5.21
# ifgroup add ipv4-group interface 192.4.5.22
# ifgroup add ipv4-group interface 192.4.5.23
# ifgroup add ipv4-group interface 192.4.5.24

# ifgroup add ipv6-group interface 2000::af:21
# ifgroup add ipv6-group interface 2000::af:22
# ifgroup add ipv6-group interface 2000::af:23
# ifgroup add ipv6-group interface 2000::af:24

# ifgroup add ipv4-group client 128.5.1.25.0/24
# ifgroup add ipv6-group client 2620::128:25:0/112
```

Note

- If no *group_name* is specified, the default group is used.
- IPv6 addresses can be entered using upper- or lower-case characters and with multiple zeroes.
- These commands properly detect any mismatches with IPv4 or IPv6 interfaces.

-
3. Select one interface on the Data Domain system to register with the backup application. It is recommended that you create a failover aggregated interface and register that interface with the backup application.

Note

It is not mandatory to choose an interface from the interface group to register with the backup application. An interface that is not part of the interface group can also be used to register with the backup application. EMC recommends that the interface be registered with a resolvable name using DNS or any other name resolution mechanism.

Note

The interface registered with the backup application is used by the backup application and its DD Boost libraries to communicate with the Data Domain system. If this interface is not available, then backups to that Data Domain system are not possible.

-
4. Once an interface and client are configured, the group is automatically enabled. Check the status (enabled or disabled) of the interface group:

```
# ifgroup status [group_name]
Status of ifgroup "default" is "enabled"
```

Note

If no *group_name* is specified, the default group is used.

5. Verify the entire configuration of all the groups with interfaces and clients:

```
# ifgroup show config all
```

Results

Sample output is displayed in the following table.

Group Name	Status	Interfaces Count	Clients Count
default	enabled	2	1
external	enabled	2	1
lab10G	enabled	2	2

Group Name	Status	Interfaces
default	enabled	10.6.109.141
default	enabled	10.6.109.41
external	enabled	10.6.109.140
external	enabled	10.6.109.142
lab10G	enabled	192.168.1.220
lab10G	enabled	192.168.1.221

Group Name	Status	Clients
default	enabled	*
external	enabled	*.datadomain.com
lab10G	enabled	ddboost-dl.datadomain.com
lab10G	enabled	yellowmedia.datadomain.com

Note

Exact name matches are done first, followed by partial name matches. So, in the example above, `ddboost-dl.datadomain` is found in the `lab10G` group.

Modifying an Interface Group

After the interface group is set up, you can add or delete interfaces from the group. The following example shows how to remove an interface from the configured interface group on the Data Domain system.

Procedure

1. Make sure that no jobs are active from the backup application to the Data Domain system on the interface you are removing from the group. You can do this from the Data Domain system by checking the status of existing connections in the interface group by enter the following command:

```
# ddbboost show connections
```

Note

Refer to documentation about the DD860 Extended Retention system (formerly the *DD860 Archiver Administration Guide*) for an example of this command on an active tier.

2. Delete an interface or client from group-name or default group on the Data Domain system.

```
# ifgroup del default interface 10.6.109.144
```

After this, the interface is released from the group and would no longer be used by the DD Boost Storage Server for any jobs from the backup servers.

Note

Removing the interface registered with the backup application makes the Data Domain system inaccessible to the backup servers. The configuration of the interface group on the Data Domain system is not deleted.

Results

To make any changes to any interface that is added to the interface group on the Data Domain system at the network layer, remove the interface from the group and add it back.

Note

If you make changes using the `net` command that modifies the interfaces, such as enabling an interface that is configured for interface group, execute the `ddboost show connections` command to update the load-balancing view. Updating the load balancing view allows the interface group to use the interface.

Removing an Interface Group

The following example illustrates removing a configured interface group on the Data Domain system.

Procedure

1. Make sure that no jobs are active from the backup application to the Data Domain system. Check the status of connections in the interface group by using the following command on a Data Domain system:

```
# ifgroup show connections
```

2. Ensure there are no pending jobs from backup servers connected to the Data Domain system.
3. Disable the *group-name* or default group on the system:

```
# ifgroup disable <group-name>
```

4. Reset the interface group:

```
# ifgroup reset <group-name>
```

Results

All the interfaces are released from the group. However, backup servers can still access the DD Boost storage server on the Data Domain system on the interface registered with the backup application.

When a group is no longer needed, use the `destroy` option to remove the group from the configuration:

```
# ifgroup destroy group-name
```

Example:

```
# ifgroup destroy external
```

Clients are matched to a group by their hostname independent of the group status (enabled/disabled). Therefore, disabling a group will not force a client to use a different group. When a client is found in a disabled group, it will use the registered interface and stay on the original connection.

Note

You can also manage interface groups from the System Manager Data Management DD Boost view. (See the *EMC Data Domain Operating System Administration Guide*).

Using Dynamic Interface Groups for MFR

The Dynamic Interface Group (DIG) feature provides the ability to control the interfaces used for DD Boost MFR to direct the replication connection over a specific network, and to use multiple network interfaces with high bandwidth and reliability for failover conditions. All Data Domain IP types are supported—IPv4 or IPv6, Alias IP/VLAN IP, and LACP/failover aggregation.

Note

The DIG feature is supported for DD Boost Managed File Replication (MFR) only.

Without the use of interface groups, configuration for replication requires several steps:

1. Adding an entry in the `/etc/hosts` file on the source Data Domain system for the target Data Domain system and hard coding one of the private LAN network interfaces as the destination IP address.
2. Adding a route on the source Data Domain system to the target Data Domain system specifying a physical or virtual port on the source Data Domain system to the remote destination IP address.
3. Configuring LACP through the network on all switches between the Data Domain systems for load balancing and failover.
4. Requiring different applications to use different names for the target Data Domain system to avoid naming conflicts in the `/etc/hosts` file.

Using interface groups for replication simplifies this configuration through the use of the DD OS System Manager or DD OS CLI commands. Using interface groups to configure the replication path lets you:

- Redirect a hostname-resolved IP address away from the public network, using another private Data Domain system IP address.
- Identify an interface group based on configured selection criteria, providing a single interface group where all the interfaces are reachable from the target Data Domain system.
- Select a private network interface from a list of interfaces belonging to a group, ensuring that the interface is healthy.
- Provide load balancing across multiple Data Domain interfaces within the same private network.
- Provide a failover interface for recovery for the interfaces of the interface group.
- Provide host failover if configured on the source Data Domain system.
- Use Network Address Translation (NAT)

The selection order for determining an interface group match for file replication is:

1. Local MTree (storage-unit) path and a specific remote Data Domain hostname
2. Local MTree (storage-unit) path with any remote Data Domain hostname
3. Any MTree (storage-unit) path with a specific Data Domain hostname

The same MTree can appear in multiple interface groups only if it has a different Data Domain hostname. The same Data Domain hostname can appear in multiple interface groups only if it has a different MTree path. The remote hostname is expected to be an FQDN, such as `dd890-1.emc.com`.

The interface group selection is performed locally on both the source Data Domain system and the target Data Domain system, independent of each other. For a WAN replication network, only the remote interface group needs to be configured since the source IP address corresponds to the gateway for the remote IP address.

Replication over LANs

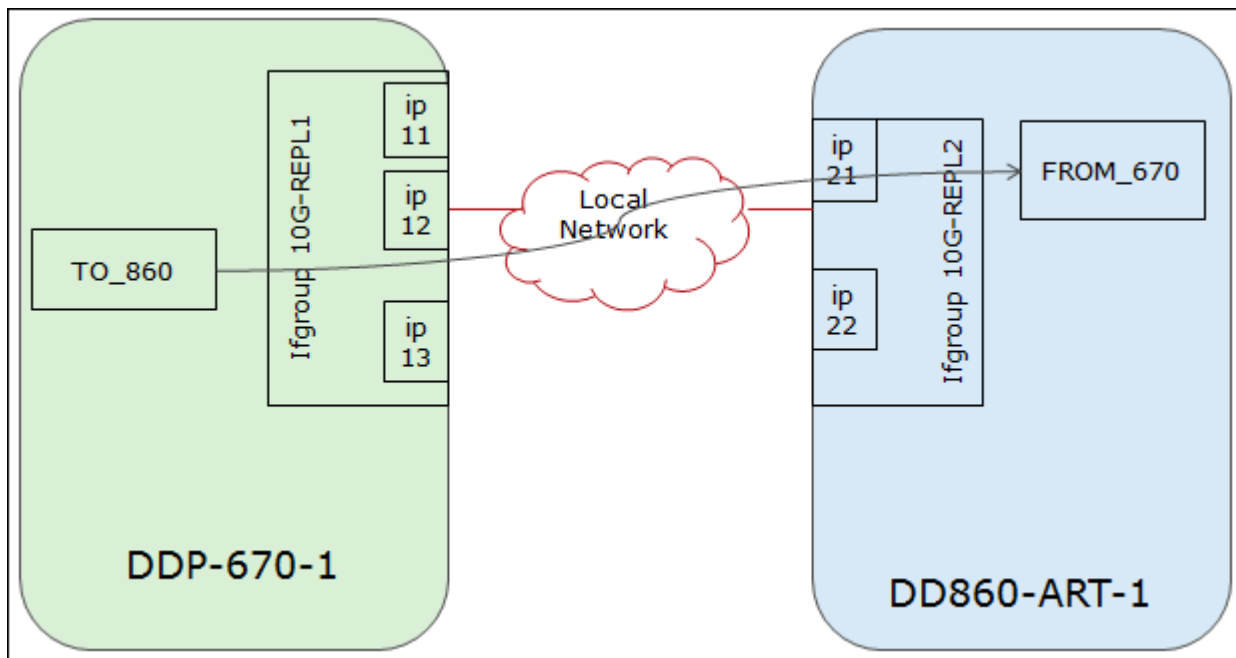
To configure interface groups for replication over a LAN:

1. Create a replication interface group on the source Data Domain system.
2. Assign for replication the local MTree path on the source Data Domain and the hostname of the destination Data Domain.
3. Create a replication interface group on the destination Data Domain system.
4. Assign for replication the local MTree path on the destination Data Domain and the hostname of the source Data Domain.

In [Figure 10](#) on page 52:

- Data Domain system DDP-670-1 is the replication source.
- Interface group 10G-REPL1 is configured with three interfaces, each of which is reachable from the destination.
- The full path to the local MTree is /data/col1/TO_860.
- DD860-ART-1 is the hostname of the remote Data Domain system (from the perspective of DDP-670-1).
- Data Domain system DDP-860-ART-1 is the replication destination.
- The replication interface group 10G-REPL2 is configured with two interfaces, each of which is reachable from the source.
- The full path to the local MTree is /data/col1/FROM_670.
- DDP-670-1 is the hostname of the remote Data Domain system (from the perspective of DD860-ART-1).

Figure 10 Using Interface Groups for Replication over a LAN



To configure the replication scenario illustrated in [Figure 10](#) on page 52:

1. Create an interface group 10G-REPL1 with three interfaces for replication on DDP-670-1. To confirm:

```
# ifgroup show config 10G-REPL1 interfaces
```

Group-name	Status	Interfaces
------------	--------	------------

```

-----
10G-REPL1      enabled  172.29.0.11
10G-REPL1      enabled  172.29.0.12
10G-REPL1      enabled  172.29.0.13
-----

```

2. Assign the full MTree path and remote hostname for replication:

```

# ifgroup replication assign 10G-REPL1 mtree /data/col1/TO_860
remote dd860-art-1
Assigned replication mtree "/data/col1/TO_860" with remote "dd860-
art-1" to ifgroup "10G-REPL1".

```

To confirm:

```

# ifgroup show config 10G-REPL1 replication

```

Group-name	Status	Replication Mtree	Replication Remote Host
10G-REPL1	enabled	/data/col1/TO_860	dd860-art-1

3. Create an interface group 10G-REPL2 with two interfaces for replication on DD860-ART-1. To confirm:

```

# ifgroup show config 10G-REPL2 interfaces

```

Group-name	Status	Interfaces
10G-REPL2	enabled	172.29.0.21
10G-REPL2	enabled	172.29.0.22

4. Assign the full MTree path and remote hostname for replication:

```

# ifgroup replication assign 10G-REPL2 mtree /data/col1/FROM_670
remote ddp-670-1
Assigned replication mtree "/data/col1/FROM_670" with remote
"ddp-670-1" to ifgroup "10G-REPL2".

```

To confirm:

```

# ifgroup show config 10G-REPL2 replication

```

Group-name	Status	Replication Mtree	Replication Remote Host
10G-REPL2	enabled	/data/col1/FROM_670	ddp-670-1

Replication over WANs

There are two options for configuring interface groups for replication over WANs:

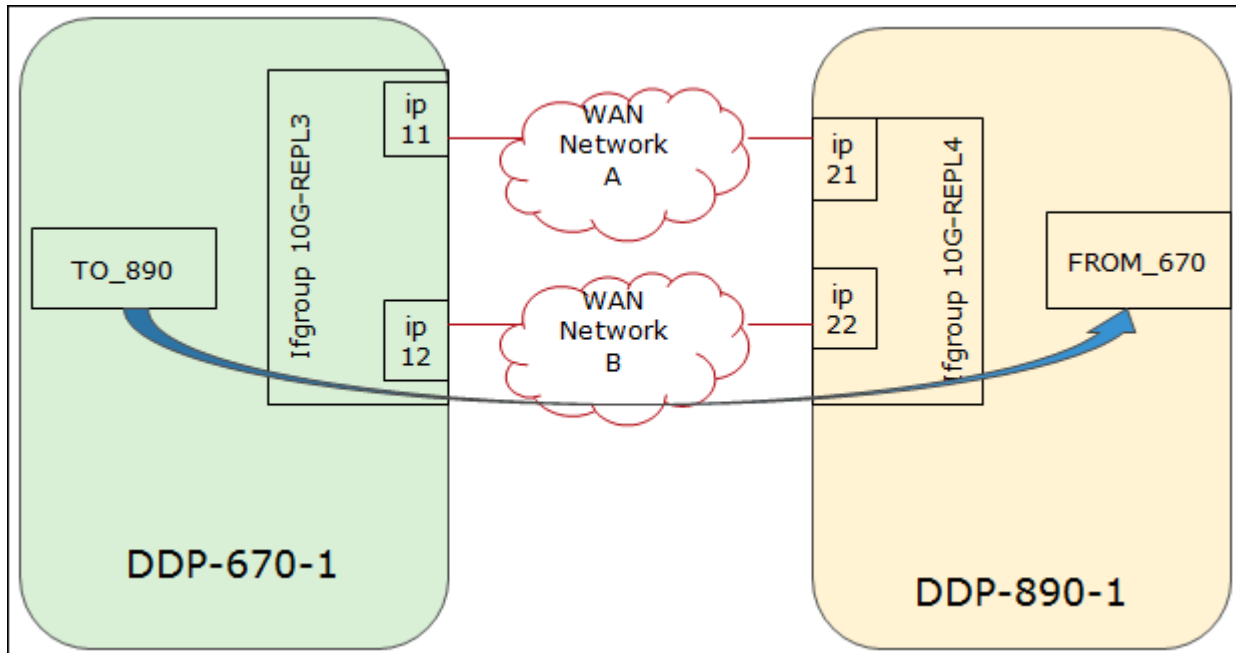
- Use the interface group on the destination only and let the source Data Domain system select the source interface based on the route configuration.
- Use interface groups on both the source and destination systems, and make sure that the IP on the source can reach the IP on the destination—you can verify that by using a trace route from both the source and destination.

Figure 11 on page 54 shows how interface groups can be used for replication over WANs:

- Data Domain system DDP-670-1 is the replication source.
- Interface group 10G-REPL3 is configured with two interfaces, each of which is reachable from one destination IP address.
- The full path to the local MTree is /data/col1/TO_890.
- DDP-890-1 is the hostname of the remote Data Domain system (from the perspective of DDP-670-1).

- Data Domain system DDP-890-1 is the replication destination.
- Interface group 10G-REPL4 is configured with two interfaces, each of which is reachable from one source IP address.
- The full path to the local MTree is /data/col1/FROM_670.
- DDP-670-1 is the hostname of the remote Data Domain system (from the perspective of DDP-890-1).
- There are two networks—WAN Network A and WAN Network B—in the same interface group.

Figure 11 Using Interface Groups for Replication over WANs




Other Supported Use Cases

Backup applications need to replicate data over their own networks and use multiple network interfaces with high bandwidth and failover capability. To support these needs, interface groups provide support for various replication paths, including multiple customer network, fan-out, and cascading paths.

Multiple Customer Network

A service provider with multiple customers may require that customers be able to replicate data across their own networks. In [Figure 12](#) on page 55, one customer is using Network A for replication of MTree11 from Data Domain System 1 (DDR1) to System 2 (DDR2), as well as MTree2 from System 2 to System 1. Another customer is using Network B for replication of MTree33 from System 1 to System 2.

Figure 12 Using Interface Groups for Replication over Multiple Customer Networks

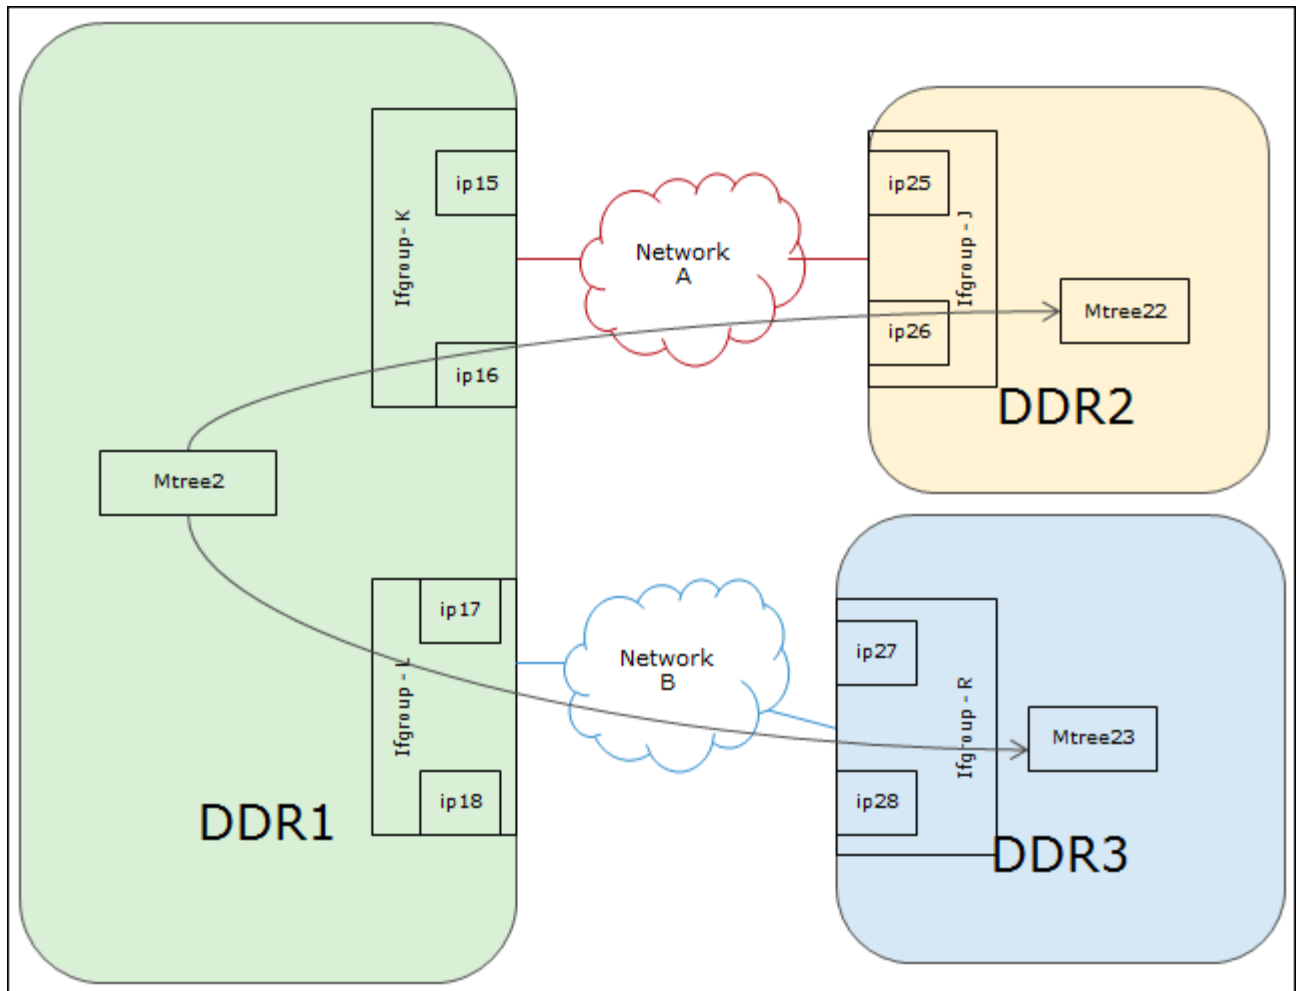


These replication paths have the same source and destination Data Domain systems, so the interface group is configured based on the MTree paths. For Network A, Ifgroup-K is configured based on the MTree11 and MTree22 paths, and Ifgroup-J is configured based on the MTree1 and MTree2 paths. For Network B, Ifgroup-L is configured based on the MTree33 path, and Ifgroup-R is configured based on the MTree3 path.

Fan-Out

In [Figure 13](#) on page 56, MTree2 is replicated from Data Domain System 1 (DDR1) to System 2 (DDR2), and from System 1 to System 3 (DDR3).

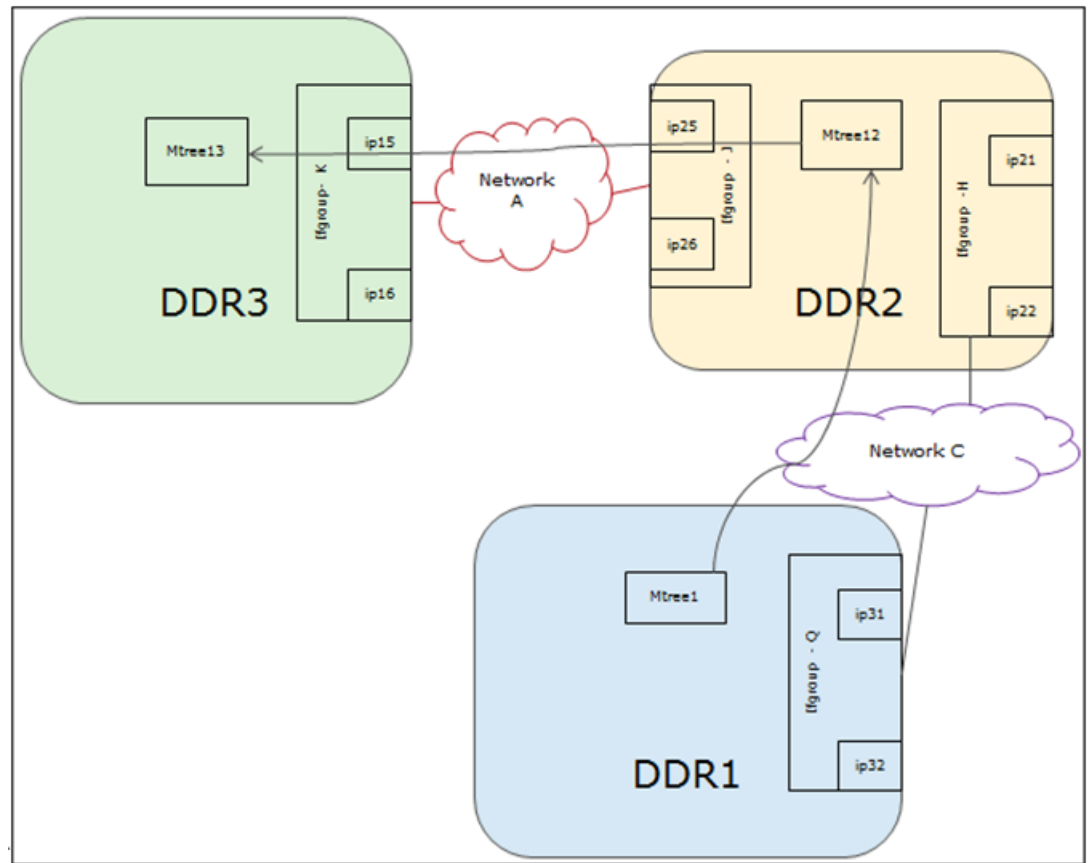
Figure 13 Using Interface Groups for Fan-Out Replication



Because the same source MTree is used, the interface group is configured based on the hostname of the remote Data Domain system. For Network A, Ifgroup-K is configured based on the hostname of Data Domain System 2, and Ifgroup-J is configured based on the hostname of Data Domain System 1. For Network B, Ifgroup-L is configured based on the hostname of Data Domain System 3, and Ifgroup-R is configured based on the hostname of Data Domain System 1.

Cascading

In [Figure 14](#) on page 57, MTree1 is replicated from Data Domain System 1 (DDR1) to System 2 (DDR2), then to System 3 (DDR3).

Figure 14 Using Interface Groups for Cascading Replication

For Network A, Ifgroup-J is configured based on the hostname of remote Data Domain System 3, and Ifgroup-K is configured based on the hostname of remote Data Domain System 2 or the path of MTree13. For Network C, Ifgroup-H is configured based on the hostname of Data Domain System 1, and Ifgroup-Q is configured based on the hostname of remote Data Domain System 2 or the path of MTree1.

Replication Failover Hostname

Supported only for DD Boost file replication, the Replication Failover Hostname feature provides the same functionality as the Failover Hostname feature for backup (see [IP Failover Hostname](#) on page 30 for details).

For replication, you configure the destination host name (remote Data Domain system) in the `/etc/hosts` file on the source Data Domain system as a failover name with another IP address that exists on the destination system.

These replication connections use the host-“failover” retry capability:

- File replication commands executed prior to the start of replication:
 - Creating a new file
 - Deleting a file
 - Finding a synthetic base file
- Requesting an interface group IP address from the destination system

Interface-group replication over IPv4 supports replication with NAT. The configuration is on the destination Data Domain system, where the source Data Domain queries for an IP address. Therefore, the destination Data Domain needs to have an interface

group with the IP addresses that the source Data Domain system can use (before NAT conversion) to reach the destination. For example:

```
Source DD -> 128.222.4.15 --NAT-> 192.169.172.8 Target DD
Source DD -> 128.222.4.16 --NAT-> 192.169.172.9 Target DD
```

On the remote Data Domain system:

```
net config ethx:1 128.222.4.15 << on the ethx of IP 192.169.172.8
net config ethy:1 128.222.4.16 << on the ethy of IP 192.169.172.8
```

Create an interface-group replication group and add the new IP addresses:

```
ifgroup create nat-repl
ifgroup add nat-repl interface 128.222.4.15
ifgroup add nat-repl interface 128.222.4.16
ifgroup replication assign remtoe <source DD hostname>
```

Network Address Translation (NAT) Support

Interface groups support Network Address Translation (NAT) for MFR. NAT is currently supported only with IPv4 addressing.

Interface group replication with NAT requires the public IP address, not the local IP address, for the destination interface group.

Note

On a Data Domain system, alias IP addresses can be configured on an active interface with IP addresses that are not reachable. In [Figure 15](#) on page 59, IP addresses 215.55.7.20 and 215.55.7.21 are aliases, configured strictly for interface group use.

To configure interface groups to support NAT:

1. Create a “natlan” interface group.
2. Create an alias IP address for the local interface.
3. Assign the external public IP address to the alias.
4. Add each new IP address to the “natlan” interface group.
5. Add replication selection (local MTree path and/or remote Data Domain hostname).
6. For the source Data Domain for replication, the source IP address of the connection needs to use the local private network.
7. For the destination Data Domain for replication, the destination IP address of the connection needs to use the public network.

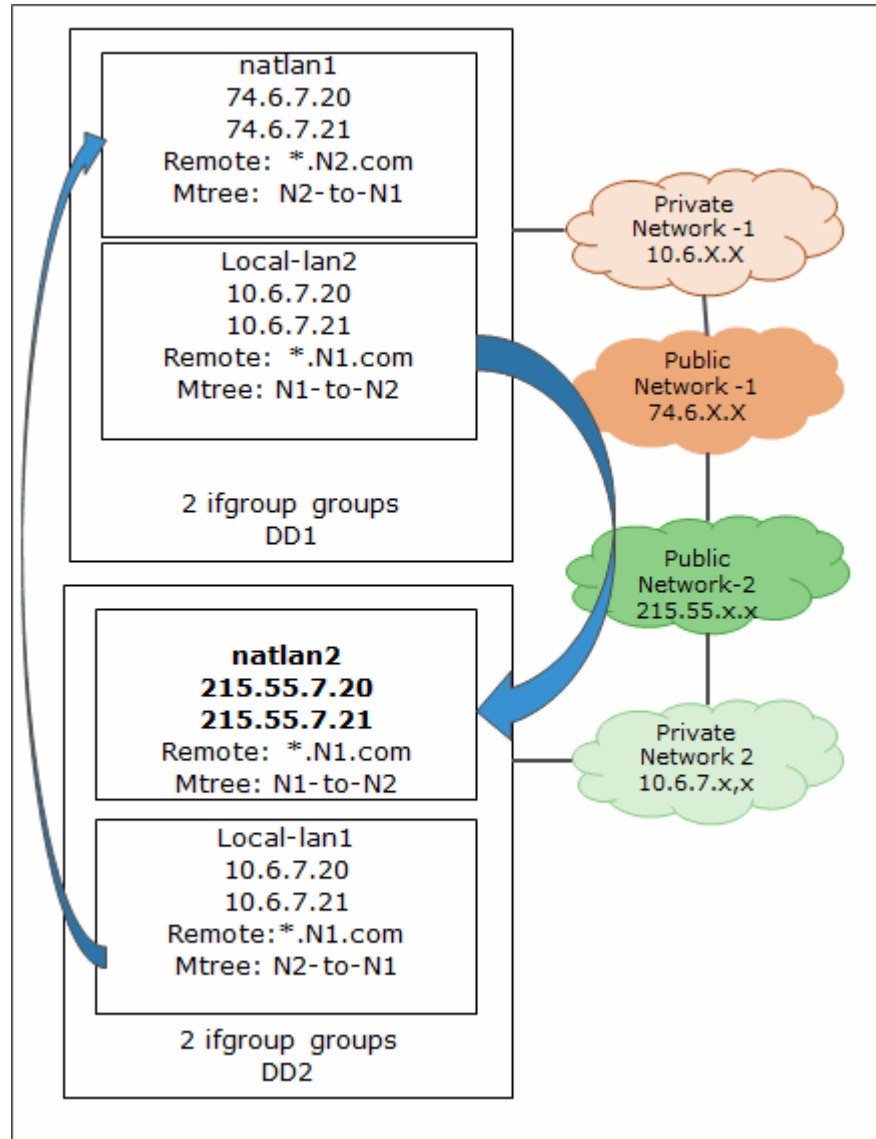
In [Figure 15](#) on page 59, DD1 and DD2 are each in a private network. Replication with NAT can be used in both directions, where the MTree name selects between the two interface groups. Only the source Data Domain system uses an interface group to select the source and destination IP. The source Data Domain queries the remote Data Domain for an IP address to use, so all IP addresses must always be from the perspective of the source Data Domain. Replication from DD1 on Network-1 to DD2 on Network-2 uses Local-lan2 with natlan2 to select both source and destination IP addresses for connection. Replication from DD2 on Network-2 to DD1 on Network-1 uses Local-lan1 with natlan1. Therefore, only the remote query for interface group IP addresses must adjust for NAT configuration.

Note

Replication with NAT requires that the following system registry entry be enabled (it is disabled by default):

```
system.IFGROUP_REPLICATION_NAT_SUPPORT = TRUE
```

Figure 15 Interface Group Replication Support for NAT



Resolving Backup/Replication Conflicts

When the registered administrative interface for backup is used for interface group replication, backup jobs erroneously select the replication interface group. To avoid this problem, add a client with the name "no-auto-detect" to the interface group that has the administrative IP address. To allow for multiple replication groups, "no-auto-detect" is internally appended with a number.

To add the client:

```
# ifgroup add repl-group client no-auto-detect
Added client "no-auto-detect" to ifgroup "repl-group".

# ifgroup show config repl-group clients

Group-name      Status      DD Boost Clients
-----
repl-group      enabled    no-auto-detect.4
-----
```

To remove the client:

```
# ifgroup del repl-group client no-auto-detect.4
Deleted client "no-auto-detect.4" from ifgroup "repl-group".
```

Configuring MFR

Enabling Low-Bandwidth Optimization

To enable the low-bandwidth option for managed file replication, enter:

```
# ddbboost file-replication option set low-bw-optim enabledLow
bandwidth optimization enabled for optimized duplication.
```

Note

Enabling or disabling the low-bandwidth optimization option does not require a restart of the Data Domain file system.

Low-bandwidth optimization can also be monitored and managed from the Enterprise Manager Data Management DD Boost view. (See the *EMC Data Domain Operating System Administration Guide*.)

No configuration changes are necessary on the backup server as this feature is transparent to the backup applications.

Note

- Enabling this feature takes additional resources (CPU and memory) on the Data Domain system, so it is recommended that this option be used only when managed file replication is being done over low-bandwidth networks with less than 6 Mbps aggregate bandwidth.
- The low-bandwidth option for managed file replication is supported only for standalone Data Domain systems.
- Low-bandwidth optimization is not supported on Data Domain Extended Retention systems.

Enabling Encryption

To enable the encrypted managed file replication option, enter:

```
# ddbboost file-replication option set encryption enabled
```

The output indicates that the encryption you requested was enabled.

No configuration changes are necessary on the backup server as this feature is transparent to the backup application. Turning on this feature takes additional resources (CPU and memory) on Data Domain system.

Enabling IPv6 Support

The existing Managed File Replication commands now include IPv4 or IPv6 functionality. For DD Boost to provide IPv6 support for managed file replication, a new keyword `ipversion` is added into the registry to provide an option to support IPv6 network. The IPv6 keyword variable is controlled through the `ddboost file-replication option set` command keyword `ipversion`. If the option `ipversion` 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.

To set the preferred IP version for DD Boost file replication to IPv6, enter:

```
# ddboost file-replication option set ipversion ipv6
Ipversion for file-replication set to "ipv6"
```

To display the current values for the DD Boost file-replication options, enter:

```
# ddboost file-replication option show ipversion
Ipversion for file-replication is: ipv6
```

To reset DD Boost file replication option to the default value IPv4, enter:

```
# ddboost file-replication option reset ipversion
```

Changing the MFR TCP Port

Changing the MFR TCP port affects all replication, not just MFR. Changing the MFR TCP port requires a restart of the Data Domain file system and should be a planned event.

To change the Replication TCP port from the default of 2051 to *port-number*, enter the following commands on both the source and destination Data Domain systems:

```
# replication option set listen-port port-number
# fileSYS restart
```

Note

Managed file replication and directory replication both use `listen-port` option. Managed file replication uses the `replication option set listen-port` command on both the source and destination to specify the port on which the destination listens and the port on which the source connects. Directory replication uses the `listen-port` option to specify only the replication destination server `listen-port`. On the replication source, the connection port for a specific destination is entered using the `replication modify` command.

- For more information on these topics, see the *EMC Data Domain Operating System Command Reference Guide*.
-

Configuring Client Access Validation

Configuring client access control for DD Boost limits access to the Data Domain system for DD Boost clients and removes dependency on the DNS. By default, if no clients are added to the clients list when DD Boost is enabled, all clients will be automatically included in the clients list. By default a `*` wildcard is used.

To restrict access, remove the `*` wildcard from the list and then add your new clients.

The backup server client list may contain both fully qualified domain names or short names. The backup host's fully qualified domain name needs to be correctly configured for reverse lookup in DNS.

To delete all clients from the DD Boost clients list, enter:

```
# ddbboost clients delete client-list
```

Optionally, to delete all clients previously added and reset the DD Boost clients list, enter:

```
# ddbboost client reset
```

Clients can be added as both fully qualified domain names and short names. To add clients to the DD Boost clients list, enter:

```
# ddbboost clients add client-list [encryption-strength {none | medium | high} authentication-mode {one-way | two-way | anonymous | kerberos}]
```

Example:

```
# ddbboost clients add ddbboost-dl.emc.com ddbboost-dl
Added "ddbboost-dl.emc.com"
Added "ddbboost-dl"
```

To view the DD Boost clients list, enter:

```
# ddbboost clients show config
```

Client	Encryption Strength	Authentication Mode
*	none	none
*.corp.emc.com	medium	anonymous
rtp-ost-ms02.domain	high	anonymous
rtp-ost-ms02.domain.com	high	anonymous

During access validation, the following search order is used to restrict access:

- Wild card * followed by partial, for example, *.emc.com followed by *.com
- Perfect match of sent client name, for example, ddbboost-dl.emc.com

If the search does not find a matching entry for the client, the client will be denied access.

Configuring DD Boost-over-FC Service

Before you begin

In order to support the DD Boost-over-FC service, it is necessary to install supported Fibre Channel Target HBAs into the system. (See also the *EMC Data Domain Operating System Command Reference Guide* and *Administration Guide* for information about `scsitarget` as a related command that may be helpful in managing the SCSI target subsystem.)

Note

- Windows, Linux, HP-UX, AIX, and Solaris client environments are supported.
 - Beginning with DD Boost 3.3 and DD OS 6.0, you do not need a device driver; you can enable DD Boost-over-FC by creating a `ddbboost fc` group; see Step 3 in the following procedure.
 - To enable DD Boost-over-FC on clients running AIX, you can also install the AIX DDdfc device driver.
-

Ensure that the client's HBA ports and the Data Domain system's endpoints are defined and that appropriate zoning has been done if you are connecting through a Fibre Channel switch.

Procedure

1. Enable the DD Boost-over-FC service:

```
# ddbboost option set fc enabled
```

2. Optional: set the DFC-server-name:

```
# ddbboost fc dfc-server-name set <server-name>
```

Alternatively, accept the default, which has the format `DFC-<base hostname>`. The hostname cannot be the fully-qualified domain name.

A valid DFC server name consists of one or more of the following characters:

- lower-case letters (“a”–“z”)
- upper-case letters (“A”–“Z”)
- digits (“0”–“9”)
- underscore (“_”)
- dash (“-”)

Note

The dot or period character (“.”) is not valid within a `dfc-server-name`; this precludes using the fully qualified domain name of a Data Domain system as its `dfc-server-name`.

Note

Similar to IP hostnames, the `dfc-server-name` is not case-sensitive. Multiple Data Domain systems accessible by the same clients using DDBoost-over-FC should be configured without case-sensitive `dfc-server-name`.

3. Create a SCSI target access group:

```
# ddbboost fc group create <group-name>
```

Example:

```
# ddbboost fc group create lab_group
```

4. To display the available list of scsitarget endpoints, enter:

```
# scsitarget endpoint show list
Endpoint          System Address      Transport           Enabled             Status
-----
endpoint-fc-0     6a                  FibreChannel       Yes                 Online
endpoint-fc-1     6b                  FibreChannel       Yes                 Online
-----
```

5. Indicate which endpoints to include in the group:

```
# ddbboost fc group add <group-name> device-set
count count endpoint endpoint-list [disk <disk-name>]
```

Note

The disk option in the previous example is optional and supported only if the client is AIX.

Example:

```
# ddbboost fc group add lab_group device-set count 8 endpoint 6a
```

6. Verify that initiators are present. To view a list of initiators seen by the Data Domain system:

```
# scsitarget initiator show list
```

7. Add initiators to the SCSI target access group:

```
# ddbboost fc group add group-name initiator initiator-spec
```

Example:

```
# ddbboost fc group add lab_group initiator
"initiator-15,initiator-16"
```

Sizing DD Boost-over FC Device-Set

The Data Domain system advertises one or more “DFC devices” of type Processor, which the DD Boost Library uses to communicate with the DD Boost-over-FC service. On the Data Domain system, access to these DFC devices is granted to one or more initiators by adding the initiator(s) to a ddbboost-type scsitarget access group:

```
# ddbboost fc group add lab_group initiator "initiator-15,initiator-16"
```

The number of DFC devices advertised to the initiator is controlled by configuring the device-set of the scsitarget access group:

```
# ddbboost fc group modify lab_group device-set count 4
```

The maximum number of supported DFC devices per Data Domain system is 64. You can have the same devices in multiple groups, but each group is limited to 64 devices.

Note

AIX DFC drivers support 128 devices; however, if you use the `disk` option with the `ddbboost fc add` command, this limitation is removed.

Because the DFC client sees each path to the Data Domain system as a separate device, more paths and more DFC devices mean better performance for constrained clients such as AIX, Windows, and Solaris.

So, how many DFC devices should be advertised to initiators on a given backup server? The answer depends upon several factors:

1. Is the backup server queue-depth constrained?

Windows platforms are considered “queue-depth constrained,” because the Windows SCSI Pass-Through Interface mechanism will only conduct 1 SCSI request at a time through each of its generic SCSI devices. This impacts the performance of the DD Boost-over FC solution, if multiple connections (e.g. backup jobs) are trying to use the same generic SCSI device. So, for Windows platforms running more than one job, it is useful to advertise multiple DFC devices.

Contrast this with the behavior of the Linux SCSI Generic driver, which imposes no such restriction. Linux is not considered “queue-depth constrained,” so it is sufficient to simply advertise one DFC device to initiators on Linux systems.

2. Number of physical paths between backup server and Data Domain system

For each advertised DFC device, the backup server operating system will create n generic SCSI devices, one for each physical path through which the backup server OS can access the device.

For example, if:

- Backup server has 2 initiator HBA ports (A and B)

- Data Domain System has 2 FC target endpoints (C and D)
- Fibre Channel Fabric zoning is configured such that both initiator HBA ports can access both FC target endpoints

then the backup server OS will see each device through four physical paths:

A -> C
 A -> D
 B -> C
 B -> D

and will create 4 generic SCSI devices for each advertised DFC device.

For a Windows backup server (with its queue-depth=1 limitation), this allows up to 4 simultaneous SCSI requests to the Data Domain system, even with only one DFC device advertised.

Sizing Calculation

The following calculation may be used to determine the number of DFC devices to advertise on the Data Domain system and to the initiators on a given backup server. EMC recommends that the same number of DFC devices be advertised to all initiators on the same backup server.

On the Data Domain System

The Data Domain system imposes a limit on the number of simultaneous requests to a single DFC SCSI device. Because of this limit, the number of devices advertised needs to be tuned depending on the maximum number of simultaneous jobs to the system at any given time. In general, the larger the number of jobs expected from backup servers using DDBoost over FC, the higher the number of devices advertised.

Let J be the maximum number of simultaneous jobs running using DFC, to the Data Domain System at any given time.

Let C be the maximum number of connections per job:

- 3 for Data Domain Extended Retention Systems
- 1 for other types Data Domain systems

Calculate:

- Maximum simultaneous connections to the DD system, using DFC, from ALL backup servers:
 - $S = J * C$
 - DFC Device Count $D = \text{minimum}(64, 2 * (S/128))$, round up
 - All DFC access groups must be configured with “D” devices.

Example:

Assume:

- 8 backup/master servers, single Data Domain systems, each server running a maximum of 50 jobs at any given time.
- Here, $J = 8 * 50 = 400$, $C = 1$ (single Data Domain system), $S = J * C = 400$, $D = 2 * 400 / 128 = 6.25$, round up to 7.
- Therefore, all DFC groups on the Data Domain system must be configured with 7 devices.

Assume:

- 8 backup servers, DD Extended Retention systems, each server running a maximum of 30 jobs at any given time.
- Here, $J = 8 * 30 = 240$, $C = 3$ (DD Extended Retention system), $S = J * C = 720$, $D = 2 * 720 / 128 = 11.25$, round up to 12.
- Therefore, all DFC groups on the DD system must be configured with 12 devices.

Linux Backup Servers

The number of DFC devices advertised on the Data Domain system using the calculations listed under [On the Data Domain System](#) is sufficient for Linux backup servers. No additional configuration is required. Linux backup servers are not queue-depth constrained, so many connections can share the same DFC generic SCSI device with no performance impact.

Windows Backup Servers

The Data Domain server path management logic spreads out connections across available logical paths (Initiator, Target Endpoint, DFC Device). We want to configure enough DFC devices such that each connection uses its own generic SCSI device (logical path) on the backup server, with a max DFC device count of 64.

Let X = the number of DFC devices configured on the Data Domain system (from [On the Data Domain System](#)). Let P = number of physical paths between backup server and Data Domain system. Let J = maximum number of simultaneous jobs, and let C = maximum number of connections per job:

– 3 for DD Extended Retention systems – 1 for other types of Data Domain systems

Calculate:

- Maximum simultaneous connections from backup server $S = J * C$, DFC device count $D = \text{minimum}((S/P), X)$, round up, up to a maximum of 64.

Note that if the value of D is greater than X , then it is sufficient to configure D devices, but only for the access group(s) with Windows clients.

Examples:

Assume:

- 4 physical paths between the backup server and Data Domain system, 30 maximum jobs, DD Extended Retention system
- In this case, $X = 25$, $P = 4$, $J = 30$, and $C = 3$
- Maximum simultaneous connections from backup server $S = (J * C) = 90$
- DFC device count $D = (90/4, 25) = 25$

So, the Data Domain system should be configured to advertise 25 devices to each initiator on the backup server.

Assume:

- 2 physical paths between the backup server and Data Domain system, 50 maximum jobs, single Data Domain system
- In this case, $X=18$, $P = 2$, $J = 40$, $C = 1$
- Maximum simultaneous connections from backup server $S = (J * C) = 40$
- DFC device count $D = \text{max}(40/2, 18) = 20$

So, the Data Domain system should be configured to advertise 20 devices to each initiator on the backup server.

Note that since the value of D (20) is greater than the value of X (18), it is sufficient to configure two devices only for the DFC access group with Windows clients.

HP-UX Backup Servers

The number of DFC devices advertised on the Data Domain system using the calculations listed under [On the Data Domain System](#) is sufficient for HP-UX backup servers. No additional configuration is required.

Note

When a Data Domain system is connected to an HP-UX host over a SAN, there is a distinct entry in the `/dev/pt` directory for each path between the Data Domain system and the HP-UX host. The entries are named `/dev/pt/pt<X>`, where `<x>` is a unique number assigned by the operating system. For example, if there are two FC ports on the Data Domain system connected to the same FC switch, and the HP-UX host has two FC ports connected to the FC switch, the entries in the `/dev/pt` directory will be `/dev/pt/pt1`, `/dev/pt/pt2`, `/dev/pt/pt3`, and `/dev/pt/pt4`.

AIX Backup Servers

For AIX, the proprietary device entries are exclusively locked on a per-process basis—one and only one process can use a device entry. Calculations are based on application instance usage. If an application spawns multiple processes, each process exclusively locks at least one device entry. Multi-threaded applications lock one device per thread. For these reasons, you should configure the Data Domain system to advertise as many DFC devices as possible (up to the maximum of 128). A `Device Busy` error may result if there are not enough devices accessible to the AIX clients.

Note

If you are using the proprietary device driver, the total number of streams in a policy should not exceed the number of AIX DFC devices available; otherwise the backup job might fail.

Solaris Backup Servers

For Solaris, device entries are exclusively locked on a per-process basis—one and only one process can use a device entry. Calculations are based on application instance usage. If an application spawns multiple processes, each process exclusively locks at least one device entry. Multi-threaded applications lock one device per thread. For these reasons, you should configure the Data Domain system to advertise as many DFC devices as possible to avoid 'in use' errors from the `sgen` device driver. A `Device Busy` error may result if there are not enough devices accessible to the Solaris clients.

The number of `sgen` devices is the number of Fibre Channel ports accessible to the Solaris instance times the number of different paths to the Data Domain system endpoint(s) times the number of LUNs in the access group.

Installing the AIX DDdfc Device Driver (Optional for AIX Clients)

DD Boost-over-FC is supported on clients running AIX versions 6.1 and 7.1. The AIX DDdfc device driver can be installed to enable the DD Boost-over-FC feature. However, you can choose instead to configure `ddboost fc group add` with the optional `disk` characteristic on the server. AIX will then treat associated DFC devices as native disk drives. The driver is packaged with the DD Boost libraries.

The driver's filename is `DDdfc.rte.1.0.0.x.bff`, where `x` is the version number. To install the driver:

Procedure

1. On the AIX client, log in as the root user.

2. Enter `# smitty install`.
3. Select **Install and Update Software**.
4. Select **Install Software**.
5. Enter the path to the `DDdfc.rte.1.0.0.x.bff` file, where *x* is the version number.
6. Press **F4** to view the list of installable items at the above path.
7. Scroll down the list until you find the `DDdfc.rte.1.0.0.x` version that you want.
8. Press **Tab** to toggle the value on the `Preview only?` line to **No**.
9. Press **Enter** to confirm your selections and install the driver.

After you finish

If you make any DFC configuration changes that impact the AIX client, execute these commands:

```
# rmdev -Rdl DDdfc
# cfgmgr
```

Do not access the DFC device(s) while these commands are executing. After these commands are run, it may take a few minutes before the configuration is effective.

If running these commands does not successfully restore full functionality, you must reboot the AIX client.

Configuring the SCSI Generic Device Driver for Solaris Clients

DD Boost-over-FC is supported on clients running Solaris 10 and 11 on SPARC and x86 hardware. DFC for Solaris uses the SCSI generic device driver (`sgen`), which is included in the Solaris installation. Use the following procedure to ensure that `sgen` successfully identifies the processor devices at startup.

Procedure

1. Add the following line in the `forceload` section of `/etc/system`:

```
forceload: drv/sgen
```

This step should resolve issues with `sgen` not properly loading during startup and keep the `sgen` driver loaded.

2. To check for existing usage of `sgen`, enter `grep sgen /etc/driver_aliases`.

Note

The existence of a `/dev/scsi`, `/dev/scsi/processor`, or `/dev/scsi/*` directory does not necessarily mean that `sgen` is currently configured. There could be dangling files.

3. If there is no existing use of `sgen`, or if `sgen` is used only for `"scsiclass,03"`, enter:

```
a. rem_drv sgen
```

```
b. add_drv -m '* 0600 root sys' -i '"scsiclass,03"' sgen
```

Note

It is critical that you use single and double quotes exactly as shown above.

This command should return to the prompt with no errors or warnings. Check connectivity to the Data Domain system. There should be at least one file in `/dev/scsi/processor`.

- c. To confirm at least one entry for three configuration files, enter: `grep sgen /etc/minor_perm /etc/name_to_major` and `/etc/driver_aliases`

Example results of this command are:

```
/etc/minor_perm:sgen * 0600 root sys
/etc/name_to_major:sgen 151
/etc/driver_aliases:sgen "scsiclass,03"
```

Note

The `name_to_major` number will likely be different than this example.

4. If the `sgen` device is already used for other devices, enter:

- a. `rem_drv sgen`
- b. `add_drv -m '* 0600 root sys' -i "scsiclass,03" "scsiclass,XX" sgen`
-

Note

XX would be the device type from a previously run `grep sgen /etc/driver_aliases`. It is critical that you use single and double quotes exactly as shown above.

An example of this command is: `add_drv -m '* 0600 root sys' -i "scsiclass,03" "scsiclass,06" sgen`.

This command should return to the prompt with no errors or warnings. Check connectivity to the Data Domain system. There should be at least one file in `/dev/scsi/processor`.

- c. To confirm at least one entry for three configuration files, enter: `grep sgen /etc/minor_perm /etc/name_to_major` and `/etc/driver_aliases`.
- d. Open the `/kernel/drv/sgen.conf` file. If the `device-type-config-list` is uncommented, add “processor” to the list to ensure that the driver is recognized. For example, if the `device-type-config-list` is uncommented as in this example:

```
device-type-config-list="direct", "sequential", "worm", "rodirect", "optical",
"changer";
```

Change the entry to:

```
device-type-config-list="direct", "sequential", "worm", "rodirect", "optical",
"changer", "processor";
```

Setting Global Authentication and Encryption

You can specify global authentication and encryption settings with DD Boost 3.4 and later.

For more information on global authentication and encryption settings, see [Authentication and Encryption Options](#) on page 17

Procedure

1. Enter the `ddboost option set` command with the type of authentication and strength of encryption you want to set:

```
ddboost option set global-authentication-mode {none | two-way  
| two-way-password} global-encryption-strength {none | medium  
| high}
```

Note

Authentication and encryption values must be set at the same time due to dependencies.

Showing Global Authentication and Encryption Settings

You can verify the current global authentication and encryption settings with DD Boost 3.4 and later.

Procedure

1. Enter the `dd boost option show` command with the arguments shown in the following example:

```
ddboost option show global-authentication-mode |  
global-encryption-strength
```

Resetting Global Authentication and Encryption Settings

You can globally reset authentication and encryption with DD Boost 3.4 and later.

Procedure

1. Enter the `dd boost option reset` command as shown in the following example:

```
ddboost option reset global-authentication-mode |  
global-encryption-strength
```

Both global values are reset to `none` when either is reset.

Note

Authentication and encryption values are reset at the same time due to dependencies.

CHAPTER 4

Backup Application Administration

Note

Complete descriptions of commands used in this guide are provided in the *EMC Data Domain Operating System Command Reference Guide*.

This chapter covers the following major topics:

- [Configuring a Backup Server](#) 72
- [Backup Administration](#) 72

Configuring a Backup Server

Configure backup servers as specified by the backup application.

Backup Administration

Network Time-Outs

Backup and restore jobs often take a long time to complete. Although the DD Boost libraries can recover from temporary network interruptions, the operating system on the backup application system might terminate a job prematurely if the backup application time-outs are set too low.

EMC recommends setting time-outs to at least 30 minutes (1800 seconds).

Note

After losing a network connection, administrators should issue the `ddboost reset stats` command to clear job connections.

CHAPTER 5

Basic Troubleshooting

This chapter provides basic troubleshooting tips that might enable customers to resolve issues on their own. For issues that cannot be resolved, customers should contact their contracted support providers.

For more information, see the Data Domain Knowledge Base, which is available at <https://support.emc.com>.

This chapter covers the following topics:

- [General Troubleshooting](#)..... 74
- [Data Domain System Settings for File Replication](#)..... 74
- [Resolve time-out error](#)..... 74
- [Managed File Replication Job Fails](#)..... 74
- [Virtual Synthetic Backup](#)..... 75

General Troubleshooting

When investigating problems, be aware that the DD Boost software has components on both a Data Domain system and a backup application system. The two environments must be compatible. The following troubleshooting considerations apply to both systems:

- **Supported Configurations**
Ensure that you have a supported configuration as specified in the *EMC Data Domain Boost Compatibility Guide* at the EMC Online Support site <https://support.emc.com>.

Note

A supported configuration can become unsupported if any component changes.

- **Authorization Failures**
If you encounter authorization failures, ensure that all of the systems have correct access credentials for the other systems. Refer to the backup application documentation for more information about access credentials.

Data Domain System Settings for File Replication

For all DD OS versions, the `replication throttle` command controls replication. Setting the throttle too low can cause file replication to fail.

Resolve time-out error

Procedure

1. Verify that the client can ping the Data Domain system.
2. Verify that the file system is running on the Data Domain system by entering:

```
# fileys status
```

3. Verify that NFS is running on the Data Domain system by entering:

```
# nfs status
```

Managed File Replication Job Fails

A typical activity monitor job detail indicates a media write error (84) occurred. The backup application log states that the NFS operation is not supported. Common causes for this error include:

- The replication license is not installed.
- Encryption is not set on both the source and destination Data Domain systems.

Add license for Replication

Procedure

1. Obtain a replication license code from Data Domain.
2. From the command-line interface on each Data Domain system, add the license code:

```
# license add license code
```

Verify Encrypted Managed File Replication Configuration

Verify the encryption option for managed file replication is enabled on both the source and destination Data Domain systems.

[Enabling Encryption](#) on page 60 describes how to enable encrypted managed file replication.

Virtual Synthetic Backup

- Verify that normal backups are OK.
- Verify that `DDP_SynWR` RPCs are being sent.

```
# ddbost show stats
```
- Verify virtual-synthetics is enabled on the Data Domain system.

```
# ddbost option show
```

