

Dell EMC Data Domain[®] Boost for OpenStorage

Version 3.4

Administration Guide

302-003-983

REV. 01

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

Published June 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	0
Chapter 1	Introducing Data Domain Boost for OpenStorage	7
	Revision History.....	8
	Overview of DD Boost for OpenStorage.....	8
	Supported Configurations	9
	Upgrade Compatibility.....	9
Chapter 2	DD Boost Features	11
	Overview of DD Boost Features.....	12
	Distributed Segment Processing	12
	In-flight Encryption.....	13
	Global Authentication and Encryption.....	14
	Methods of Setting Authentication and Encryption.....	14
	Authentication and Encryption Settings.....	14
	Authentication and Encryption Options.....	15
	Backwards Compatibility Scenarios.....	16
	Managed File Replication (MFR).....	18
	Low-Bandwidth Optimization.....	18
	Encrypted Managed File Replication.....	19
	DD Boost and High Availability.....	19
	DD Boost, HA, and failover.....	19
	Partial HA configurations	20
	Auto Image Replication (AIR).....	20
	Limitations When Using AIR on Data Domain Systems	22
	Targeted Auto Image Replication.....	22
	MTree Replication.....	25
	IPv6 Support.....	26
	Dynamic Interface Groups: DD Boost IP Data Path Management.....	26
	Interfaces.....	27
	Clients.....	30
	DD Boost-over-Fibre Channel Transport.....	32
	DD Boost-over-Fibre Channel Path Management.....	34
	Initial Path Selection.....	36
	Dynamic Re-Balancing.....	36
	Client Path Failover.....	36
	Queue-Depth Constraints.....	37
	Virtual Synthetic Backups.....	37
	Client Access Validation.....	38
	DD Boost Multiuser Data Path.....	38
	Storage Unit Management.....	38
	Multiuser Storage Units Access Control.....	38
	Storage Unit Capacity Quotas.....	39
	Storage Units Stream Count Management.....	39
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.....	51
	Replication over LANs.....	52
	Replication over WANs.....	54
	Other Supported Use Cases.....	55
	Replication Failover Hostname.....	58
	Network Address Translation (NAT) Support.....	59
	Resolving Backup/Replication Conflicts.....	60
	Configuring MFR.....	61
	Enabling Low-Bandwidth Optimization	61
	Enabling Encryption.....	61
	Enabling IPv6 Support.....	62
	Changing the MFR TCP Port.....	62
	Configuring Client Access Validation.....	62
	Enabling In-flight Encryption.....	63
	Configuring DD Boost-over-FC Service.....	64
	Sizing DD Boost-over FC Device-Set.....	65
	Sizing Calculation.....	67
	Installing the AIX DDdfc Device Driver (Optional for AIX Clients)....	69
	Configuring the SCSI Generic Device Driver for Solaris Clients.....	70
	Setting Global Authentication and Encryption.....	72
	Showing Global Authentication and Encryption Settings.....	72
	Resetting Global Authentication and Encryption Settings.....	72
Chapter 4	Installing DD Boost for OpenStorage	75
	Installation Overview.....	76
	OST Plug-In and DD OS Upgrades.....	76
	Firewalls and Ports.....	77
	Installing OST Plug-In for NetBackup.....	77
	Installing the OST Plug-In on Media Servers.....	78
	Installing the UNIX Plug-In	78
	Installing the Windows Plug-In.....	78
	NetBackup Services.....	79
	Installing OST Plug-In for Backup Exec.....	80
	Installing the Plug-In on Media Servers.....	80
	Install the Windows Plug-In.....	80
	Backup Exec Services.....	81
	Tuning Windows Media Servers for Performance.....	81
	Uninstalling the Windows Plug-in.....	81
Chapter 5	Backup Application Administration	83
	Configuring a Media Server.....	84
	NetBackup Configuration.....	84
	Backup Exec Configuration.....	94
	NetBackup Administration.....	98
	Find your OST Plug-in Version.....	98
	Find your NetBackup version.....	98

	Network Time-Outs.....	98
	Grouping Storage Units to Provide Failover.....	99
	Backup Exec Administration.....	100
	Find your OST plug-in version.....	100
	Find your Backup Exec version.....	100
	Delete Storage Units on Data Domain Systems.....	100
Chapter 6	Basic Troubleshooting	101
	General Troubleshooting.....	102
	Data Domain System Settings for File Replication.....	102
	NetBackup Troubleshooting.....	102
	Unable to Delete the Data Domain System.....	102
	Check the Installation.....	103
	Check Credentials.....	104
	Resolve License Errors.....	104
	Error Logging on the Media Servers.....	104
	Resolving Failed Backups on Media Servers.....	105
	Resolve Plug-In Log Messages.....	106
	Resolve “Cannot connect on socket” Error.....	106
	NetBackup Backup Jobs Fail on Solaris Media Servers.....	106
	Backup jobs from AIX clients fail.....	107
	Optimized Duplication Job Fails.....	107
	Virtual Synthetic Backup.....	108
	Monitoring Auto Image Replication.....	108
	Backup Exec Troubleshooting.....	113
	Basic Troubleshooting.....	113
	Check the installation.....	113
	Check Credentials for a Data Domain System.....	114
	Resolve License Errors	114
	Set Up Active Debugging.....	114

CONTENTS

CHAPTER 1

Introducing Data Domain Boost for OpenStorage

This chapter contains the following topics:

- [Revision History](#) 8
- [Overview of DD Boost for OpenStorage](#)..... 8
- [Supported Configurations](#) 9
- [Upgrade Compatibility](#)..... 9

Revision History

The following table presents the revision history of this document.

Table 2 Revision History of DD Boost for OpenStorage Administration Guide, Release 3.4

Revision	Date	Description
01 (3.4)	June 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.

Overview of DD Boost for OpenStorage

In the context of Veritas backup applications (NetBackup and Backup Exec), Data Domain Boost (DD Boost) has two components:

- An OST plug-in that you install on each media server. This plug-in includes the DD Boost libraries to integrate with the DD server that runs on the Data Domain system.
- The DD 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 5.7.1, it can also be a dual-node Data Domain HA system.

The backup application (NetBackup or Backup Exec) sets policies that control when backups and duplications occur. Administrators manage backup, duplication, and restores from a single console and can use all of the features of DD Boost, including WAN-efficient replicator software.

The Data Domain system exposes pre-made disk volumes called storage units to a DD Boost-enabled media server. Multiple media servers, each with the Data Domain OST plug-in, can use the same storage unit on a Data Domain system as a storage server. Each media server can run a different operating system, provided that the operating system is supported by Data Domain and the backup applications NetBackup or Backup Exec.

The figure shows an example configuration of Data Domain Boost for Open Storage using NetBackup.

- 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.

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](#)..... 12
- [Distributed Segment Processing](#) 12
- [In-flight Encryption](#)..... 13
- [Global Authentication and Encryption](#)..... 14
- [Managed File Replication \(MFR\)](#)..... 18
- [DD Boost and High Availability](#)..... 19
- [Auto Image Replication \(AIR\)](#)..... 20
- [MTree Replication](#)..... 25
- [IPv6 Support](#)..... 26
- [Dynamic Interface Groups: DD Boost IP Data Path Management](#)..... 26
- [DD Boost-over-Fibre Channel Transport](#)..... 32
- [DD Boost-over-Fibre Channel Path Management](#)..... 34
- [Virtual Synthetic Backups](#)..... 37
- [Client Access Validation](#)..... 38
- [DD Boost Multiuser Data Path](#)..... 38
- [Storage Unit Management](#)..... 38

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.

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).

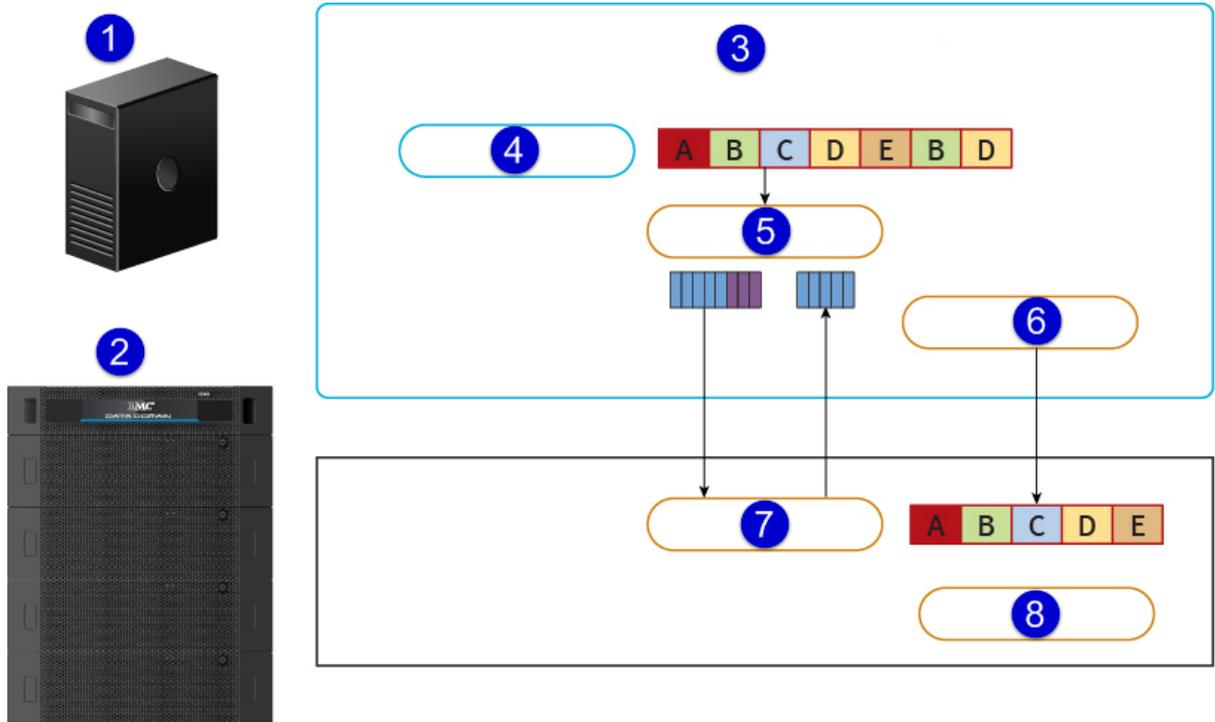
Distributed Segment Processing

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 OpenStorage Plug-in
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-+ class machines (T4, T5, M6, M7) 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.

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 two ways.

- 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 only, the client setting determines the minimum authentication and encryption. Connection attempts that do meet these settings will fail. If per-client settings are defined when global settings have been defined, the per-client values will be used only if they match or are greater than the global settings. Per-client settings that are less than the global settings are ignored.

For example:

- If a client is configured to require anonymous authentication and the global authentication setting is "two-way password," then two-way password authentication is used since that is the stronger form of authentication.
- If the client is configured with the authentication setting "two-way password" and the global setting is "none," then "two-way-password" authentication is used since that is the stricter setting.

Caller-specified values

Caller-specified values are used if they are stricter than either the global or per-client settings.

Authentication and Encryption Options

You can select one of three allowed settings for both global authentication and encryption. These settings range from the least secure (but most backwards-compatible) to the most secure.

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`.

Data Domain systems support additional per-client and global authentication values that you cannot use with DD Boost for OpenStorage. These values are "one-way" and "two-way."

Authentication settings

The following list ranks authentication values from weakest to strongest:

1. none
Not secure; this is the default setting.
2. anonymous
In-flight data is encrypted.
3. two-way password
This is the most secure option.
In-flight data is encrypted.

Encryption settings

The following list ranks encryption values from weakest to strongest:

1. none
Not secure; this is the default setting.
2. medium
Employs AES 128 and SHA-1.
3. high
Employs AES 128 and SHA-1.

Global authentication

The three `global-authentication-mode` options offer different levels of protection and backwards compatibility.

For a complete list of DD Boost commands and options, see the *Data Domain Command Reference Guide for Version 6.1*. However, additional authentication settings described in that document cannot be used with DD Boost for OpenStorage.

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.

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 or per-client settings for DD Boost OpenStorage (OST) clients on the Data Domain system must be set to "none" or "anonymous."
- 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 OST 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 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](#).

Table 3 One Setting

Call specifies	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 specifies	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.

Table 4 Multiple Settings (continued)

Call specifies	Per-client settings	Global settings	Used values
			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
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 51.

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 *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 *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 make 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.

Auto Image Replication (AIR)

Auto Image Replication (AIR) works by duplicating images to a remote master server domain. The AIR feature, introduced in NetBackup (NBU) 7.6, addresses the site to site replication challenge by allowing Storage Lifecycle Policies (SLP) to duplicate selected images between NetBackup Master domains.

The primary purpose of AIR is to create off-site copies of mission critical backups to protect against site loss. It is not intended to extend the storage capacity of a backup domain by allowing backups to be stored in a separate domain; nor is it intended to provide for day-to-day restores of data. Due to WAN bandwidth restrictions between sites, typically only the most critical data should be chosen for replication using AIR. Electronic off-siting in this manner allows the backup set to be replicated to an off-site location as soon as the backup has completed at the primary site without the need for user intervention based on the configuration of the SLP. It also means that the replicated copy is available at the disaster recovery site as soon as the duplication has completed.

To use AIR, suitable disk storage units must be configured in the source and target domains. The storage units are associated with each other using management `ddboost association` commands configured on each Data Domain system.

The figure illustrates the following configuration:

The source Data Domain system (D1) provides the routing for the backup image copies to the target domain:

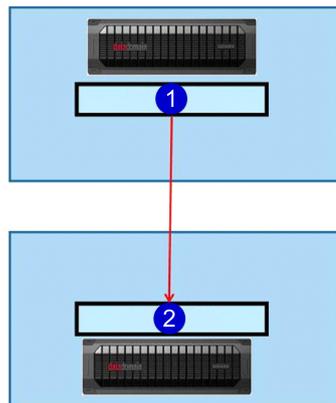
```
ddboost association create D1-SU-A replicate-to D2 D2-SU-B
```

The target Data Domain system (D2) provides for the authentication and event notification:

```
ddboost association create D2-SU-B replicate-from D1 D1-SU-A
```

Currently only one association for each storage-unit is supported. So only the single replicate scenario is supported, as shown in the following figure:

Figure 3 Auto Image Replication



1. D1 (NBU Domain 1) SU-A
2. D2 (NBU Domain 2) SU-B

Note

Only IP server names are valid in creating AIR associations, DFC server names should not be used in creating AIR associations.

Auto Image Replication works by duplicating backups from a disk pool in the source domain to a disk pool in the target domain. The replication operation requires two SLPs, one in the source domain and one in the target domain, both of which must have the same name. The SLP in the source domain is associated with the backup policy and controls the writing of the backup and the subsequent replication to the target domain.

The SLP in the target domain is not associated with a backup policy but is invoked by an alerting mechanism (an import) when a new image duplicated from the source domain is detected. This SLP runs the process to add the information about the backup to the target domain and can also be configured to replicate the backup to other storage locations in the target domain. The master media server (MMS) in the remote NBU domain will create the importing SLP by default should the user choose not to setup the remote MMS SLP.

When the backup to the primary domain is complete, the backup is replicated to the destination domain and the catalog at the destination domain is updated. This gives the backup operator at the target domain the ability to restore data to clients in the target domain in the event of a catastrophic loss of the entire source NBU domain environment.

Limitations When Using AIR on Data Domain Systems

- AIR is only supported to a single-target storage unit. Replications cannot be cascaded using a single SLP. Replications can be cascaded from the originating domain to multiple domains if an SLP is set up in each intermediate domain to anticipate the originating image, import it, and then replicate it to the next target master.
Cascades may not result in a circular loop-back condition, direct or indirect, to the original source storage unit.
- For NetBackup customers using the AIR feature with a large number of policies associated with a single LSU, a feature is available that may improve performance. For more information, contact Online Support at <https://support.emc.com>.

For additional AIR limitations, known issues, and workarounds, please refer to the *Data Domain Boost for OpenStorage Release Notes*.

Targeted Auto Image Replication

Targeted Auto Image Replication (AIR) allows Data Domain systems running DD Boost 3.3 and NetBackup 7.7.3 to replicate a backup job to multiple destination storage units. Targeted AIR also allows multiple storage units to replicate backup jobs to a single destination, in addition to other configurations.

AIR provides OST plug-in support for the NetBackup application to replicate an image set that has completed backup. Targeted AIR builds on the existing AIR functionality, allowing you to use `ddboost association` commands to define one-to-many, many-to-one, many-to-many, and cascade associations.

Targeted AIR has the following software requirements:

- DD OS version 6.0 installed on all Data Domain systems
- DD Boost OST plug-in 3.3 installed on all media servers
- NetBackup 7.7.3 installed on all media servers

Targeted AIR does not require the following:

- Any special hardware beyond the previous AIR requirements
- An additional license if you are already using AIR

Supported configurations for Targeted AIR

Targeted AIR allows you to replicate a backup job to multiple destination storage units. It also allows multiple storage units to replicate backup jobs to a single destination.

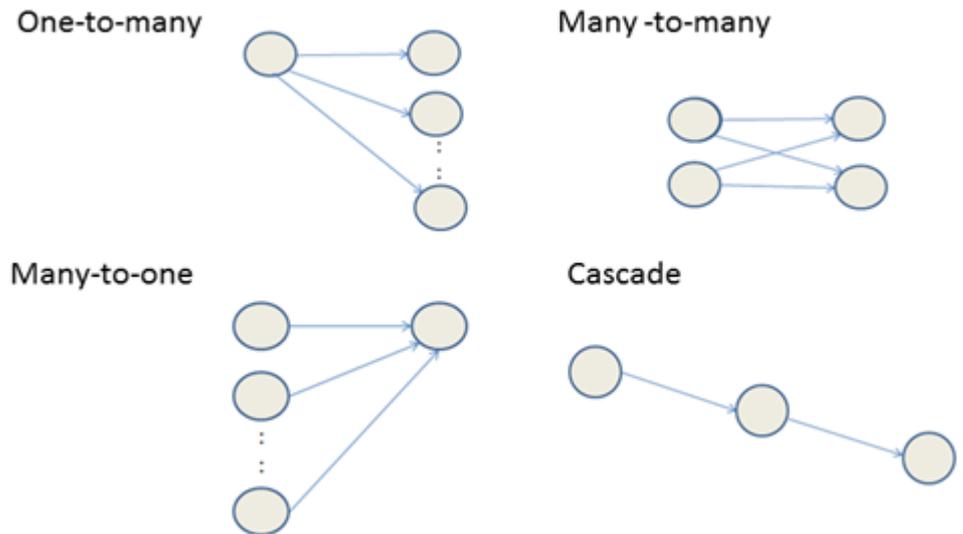
Several topologies support Targeted AIR:

- **One-to-many**
A Targeted AIR configuration in which an image set in a source storage unit can be replicated to multiple target storage units.
- **Many-to-one**
In this configuration, image sets in multiple source storage units can be replicated to a single storage destination.
- **Many-to-many**
In this configuration, image sets in multiple source storage units can be replicated to multiple storage destinations.
- **Cascade**

A replication configuration in which replication moves from Storage Unit A to Storage Unit B to Storage Unit C. A cascade topology does not allow a circular loopback configuration; Storage Unit C does not replicate back to Storage Unit A.

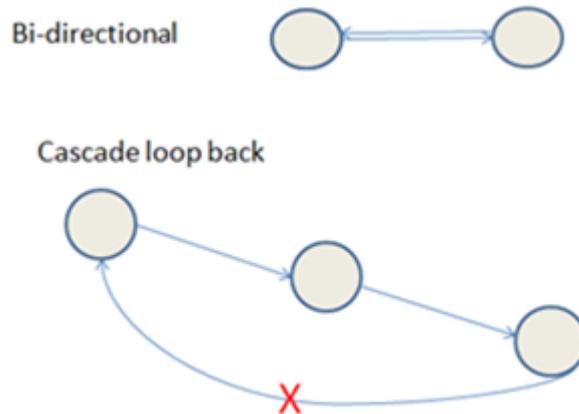
In the following illustration, the circles represent storage units, and the connecting lines represent the DD Boost associations and data communication links:

Figure 4 Supported Targeted AIR DD Boost associations



Other topologies do not support Targeted AIR. With bidirectional replications, the potential jobs could run simultaneously in both directions using the same filename; there is no advantage in this case because the source and target storage unit are the same, resulting in the same backup images and their replications being in the same storage unit.

Figure 5 Topologies that do NOT support Targeted AIR DD Boost associations



Association Tables

The `ddboost association` commands use the Local Storage Unit as the table search key.

The following is an example of the `ddboost association show` report:

```
-----
# ddboost association show
Local Storage Unit  Direction          Remote Host      Remote Storage Unit  Import To
-----
SU_SRC_A           replicate-to      localhost       SU_TRG_A             -
SU_TRG_B           replicate-from    localhost       SU_SRC_B             -
-----
```

The association table is stored in the DD OS registry. For the previous association table report, the DD OS registry is as follows:

```
# reg show protocol.ost.air.association
protocol.ost.air.association.SU_SRC_A.1 = replicate-to,localhost,SU_TRG_A
protocol.ost.air.association.SU_TRG_B.1 = replicate-from,localhost,SU_SRC_B,(null)
```

The `null` that follows the second registry entry, `SU_TRG_B`, indicates there is no “Import To” association defined. This corresponds to a dash - that appears in the `ddboost association show` report for Local Storage Unit `SU_TRG_B`.

Although associations can be displayed using the command `reg show`, only the `ddboost association` commands should be used to change association table entries.

When you to check the association table, make the `ddboost association show` command is used in the context of the Local Storage Unit of the Data Domain system where the command is executed.

The `ddboost association create` command checks to make sure the number of associations defined for the Local Storage Unit does not exceed the limits defined for either `system.REPLICATE_MAX_ASSOCIATIONS` or

`system.REPLICATE_MAX_TO_ASSOCIATIONS`. If either limit is exceeded, the command fails and you see an error message similar to the following:

```
When SE@admin## ddbost association create SU_TRGB replicate-from localhost TEST_LSU_7
Number of associations for SU_TRGB exceeded the maximum allowed, [6]
SE@admin## ddbost association create SU_SRC replicate-to localhost SU_TRGB_7
Number of associations for SU_SRC exceeded the maximum allowed, [6]
```

Specifying Targeted AIR for different configurations

When you want to configure Targeted AIR to replicate backup jobs to targets, you need to specify each destination system using the `ddbost association` command.

You configure each target system as if the association is one-to-one, repeating the commands for each system that receives a copy of the backup image.

Procedure

1. Enter the `ddbost association create` command to enable the source to replicate to a new target.

In the following example, the source Data Domain system (D1) provides the routing for the backup image copies to the target domain:

```
ddbost association create D1-SU-A replicate-to D2 D2-SU-B
```

The target Data Domain system (D2) provides for the authentication and event notification:

```
ddbost association create D2-SU-B replicate-from D1 D1-SU-A
```

2. If you want to add an additional target system, (D3), repeat Step 1 for D3 to complete the configuration.

```
ddbost association create D1-SU-A replicate-to D3 D3-SU-B
```

As in the previous example, the additional target system (D3) provides for the authentication and event notification.

```
ddbost association create D3-SU-B replicate-from D1 D1-SU-A
```

3. Repeat the previous commands for as many supported configurations you wish to create.

See [Supported configurations for Targeted AIR](#) on page 22 for more information.

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 `ddbost storage-unit modify` command. (See the *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 *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 51.

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 30.

Interfaces

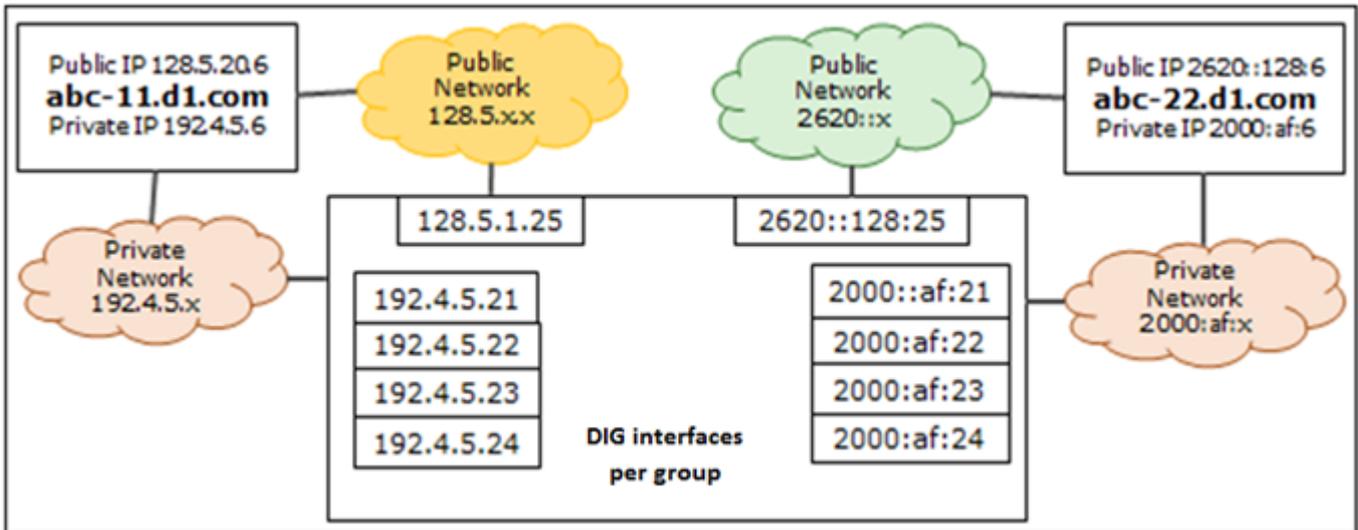
A Dynamic Interface Group (DIG) interface is a member of a single interface group *<group-name>* and may have the following characteristics:

- 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 6 DIG Support for IPv4 and IPv6 Addressing



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"`.

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
```

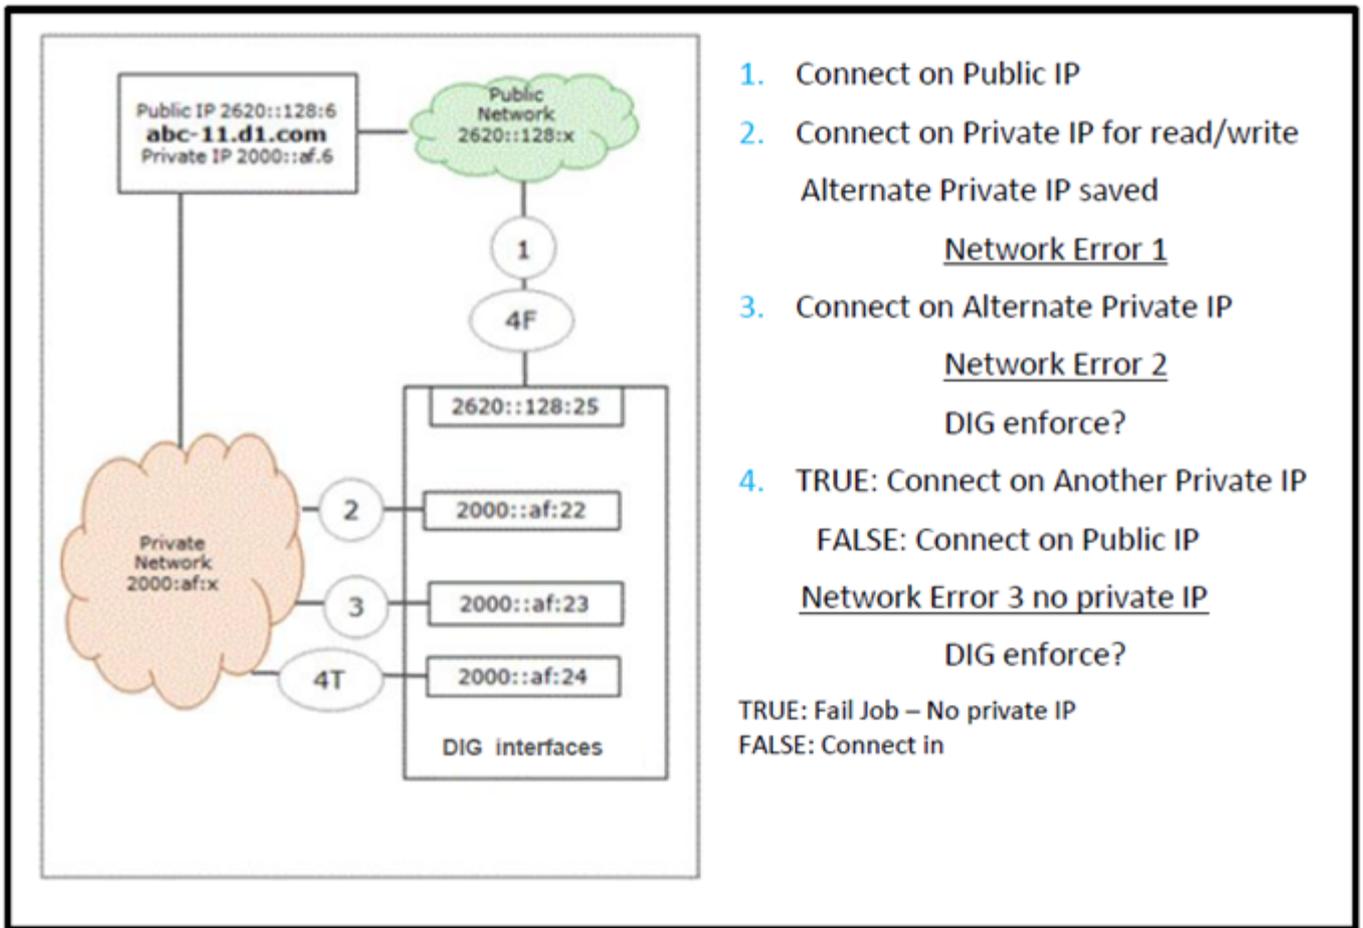
Note

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 7 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 8](#) on page 31):

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 8](#) on page 31) 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`

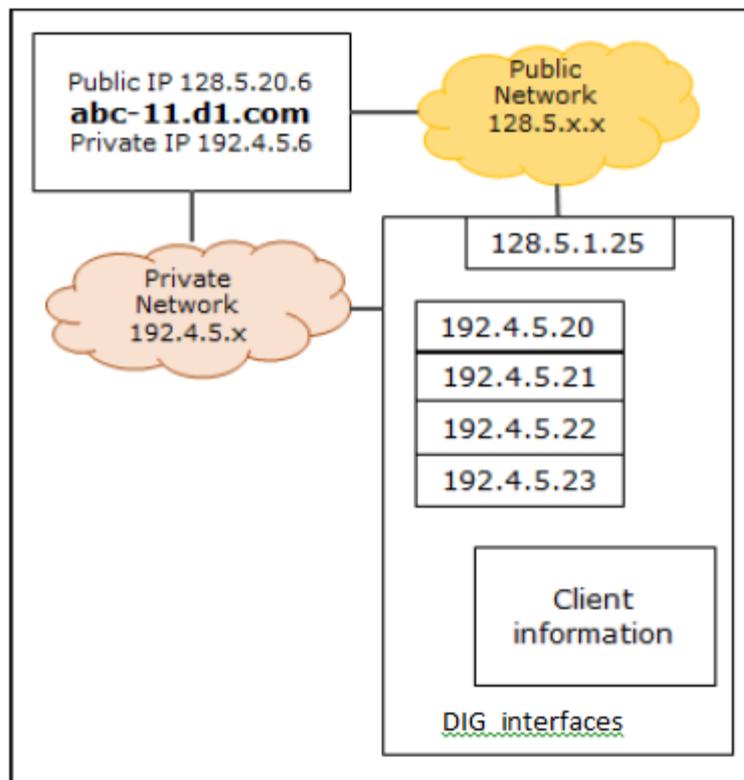
5. All Clients: `*`

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.

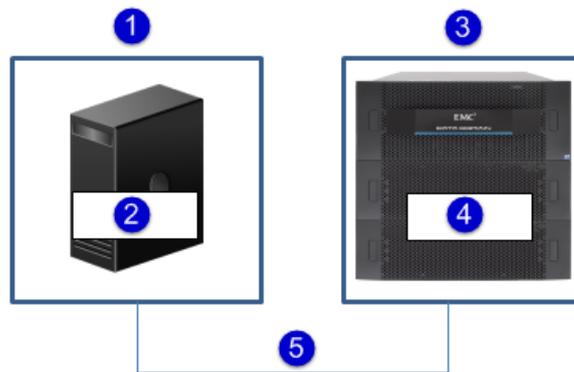
Figure 8 DIG Host Range for Client Selection



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 9](#) on page 32.

Figure 9 DD Boost-over-IP Transport



1. Media 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

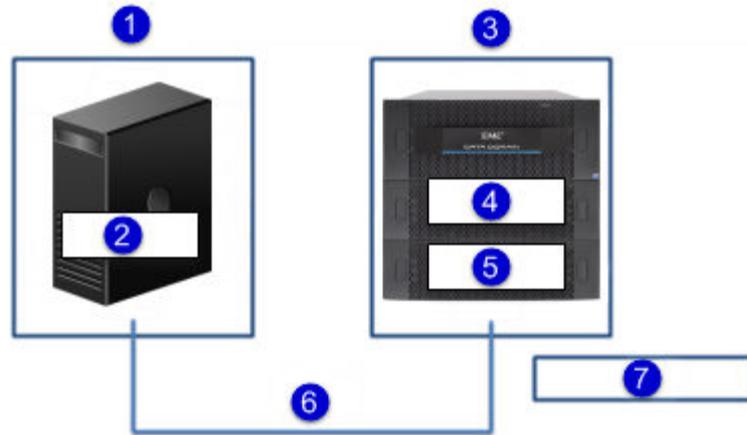
Windows, Linux, HP-UX, AIX, and Solaris client environments are supported.

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 10 SCSI Commands between Media Server and Data Domain system.



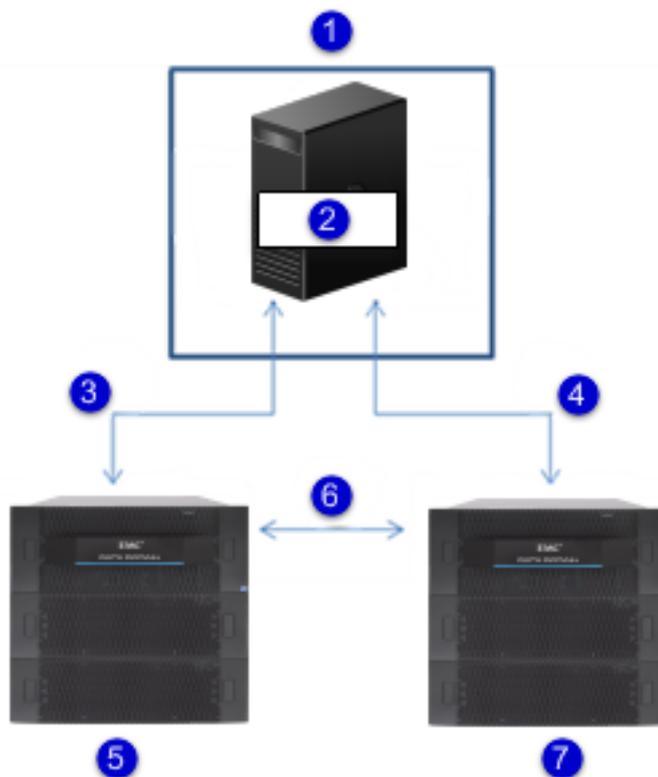
1. Media Server
2. Application, DD Boost Library, DD Boost-over-FC Transport
3. Data Domain System
4. DD Boost Service
5. DD Boost-over-FC Server
6. SCSI Commands over FC
7. SCSI Processor Devices

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 media 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 media server and the Data Domain systems.

Figure 11 Fibre Channel Communication Path

1. Media 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.

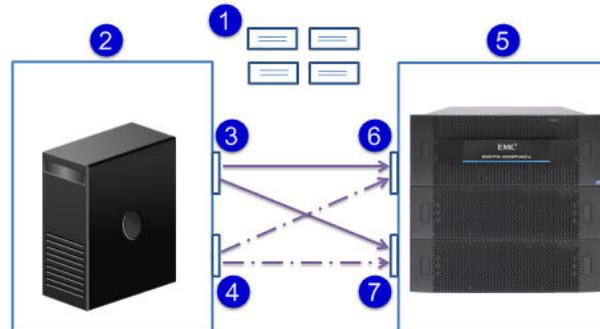
The Data Domain system advertises one or more SCSI processor-type devices to the media server, over one or more physical paths. The operating system discovers all devices through all available paths, and creates a generic SCSI device for each discovered device and path.

For example, consider the case where:

- Media 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 12 DD Boost-over-FC Path Management Scenario



1. Four Devices
2. Media 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 media 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 65 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 media 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 media 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.

Data Domain also supports “in-line” virtual synthetic backups in which each subsequent incremental backup is merged in-line with the current full backup, producing a full backup image set on the Data Domain system. After each incremental backup, you have a full backup image set ready for a restore operation. This ensures you always have the latest full backup and do not need to specifically run the full synthetic backup to “stitch together” the base and all the incremental backups that have accumulated to that point.

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 `ddboost` 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 (NetBackup or Backup Exec) 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 `ddboost 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 *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 `ddboost storage-unit modify`. You can display the reported physical size using `ddboost storage-unit show`.

See the `ddboost`, `quota`, and `mtree` sections of the *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 (such as Veritas NetBackup and Backup Exec) 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 `ddboost storage-unit create` command. To configure stream limits for an existing storage unit, use the `ddboost 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

The following procedures for configuring a Data Domain system apply to NetBackup and Backup Exec.

Note

Complete descriptions of commands used in this guide are provided in the *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](#)..... 51
- [Configuring MFR](#)..... 61
- [Configuring Client Access Validation](#).....62
- [Configuring DD Boost-over-FC Service](#)..... 64
- [Setting Global Authentication and Encryption](#)..... 72

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 provided license key.

Refer to the applicable *Data Domain Operating System Release Notes* for the most up-to-date information on licensing and service.

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 *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:

```
# ddboost 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:

```
# ddboost user show
```

DD Boost user	Default tenant-unit	Using Token Access
ddbu1	Unknown	Yes
ddbu2	Unknown	-
ddbu3	Unknown	Yes
ddbu4	Unknown	-
ddbu5	Unknown	-
ddbu6	Unknown	-
ddbu7	Unknown	Yes
ddbu8	Unknown	-

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

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

Note

DD Boost Token Access will not support upgrades across major releases.

Creating Storage Units

You need to create one or more storage units on each Data Domain system enabled for OpenStorage in a NetBackup or Backup Exec installation. In a NetBackup system, a storage unit corresponds to disk pools on the media server whereas in a Backup Exec system, it corresponds to a tape repository.

Procedure

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

```
# ddboost 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 *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:

```
# ddbost 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:

```
# ddbost 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 *Data Domain Operating System Command Reference Guide*.

Note

If quotas are enabled, some OpenStorage backup applications may report unintuitive sizes and capacities. A Knowledge Base article, “Storage Unit Quota Display on NetBackup and Backup Exec” (Document ID 000185210), has been developed to explain this in more detail.

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 the hard stream limit is exceeded, NetBackup and Backup Exec report the error to the client and automatically retry read and write jobs.

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 media servers and the OST plug-ins 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 OST plug-in 2.2 or later 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.

It is a best practice to register the interface with a resolvable name using DNS or any other name resolution mechanism. Using NetBackup and assuming that 192.168.1.1 is named `dd22.abc.com`, execute the following command on the media server:

```
nbdevconfig -creatests -st 9 -stype DataDomain -storage_server
dd22.abc.com -media_server load64
```

Note

The interface registered with the backup application is used by the backup application and its OST plug-in 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 media servers.

Note

Removing the interface registered with the backup application makes the Data Domain system inaccessible to the media 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 media 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, media 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 *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.domain.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.

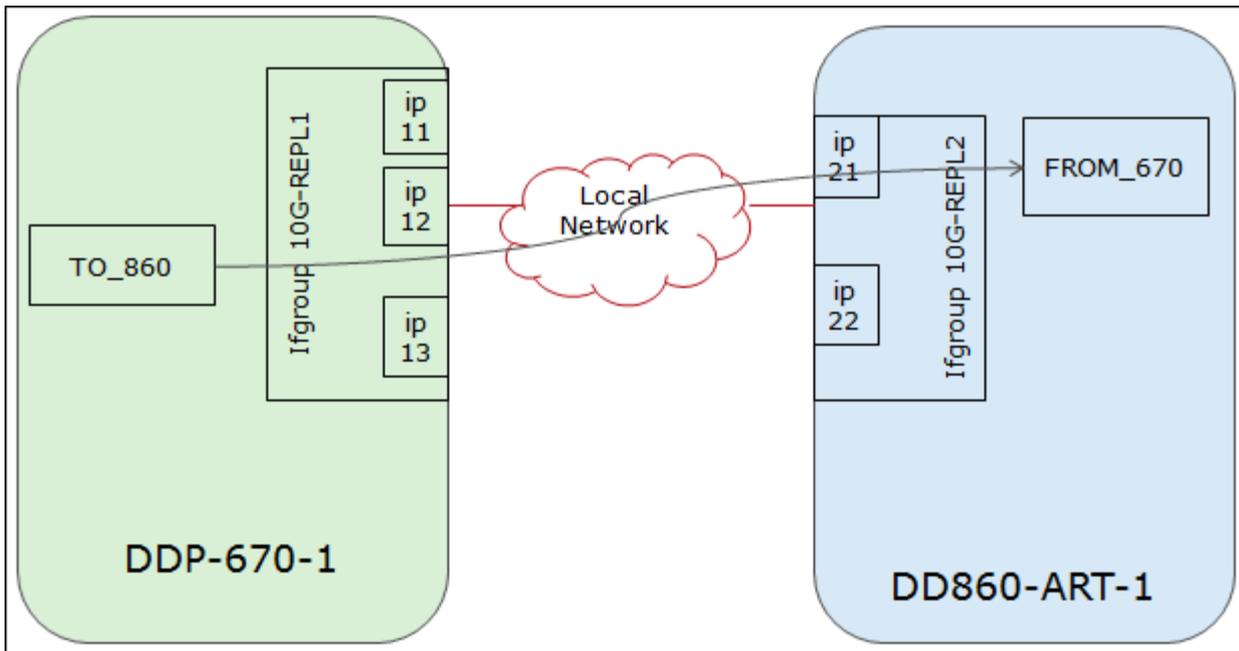
Replication over LANs

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 13](#) on page 53:

- 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 13 Using Interface Groups for Replication over a LAN

To configure the replication scenario illustrated in [Figure 13](#) on page 53:

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

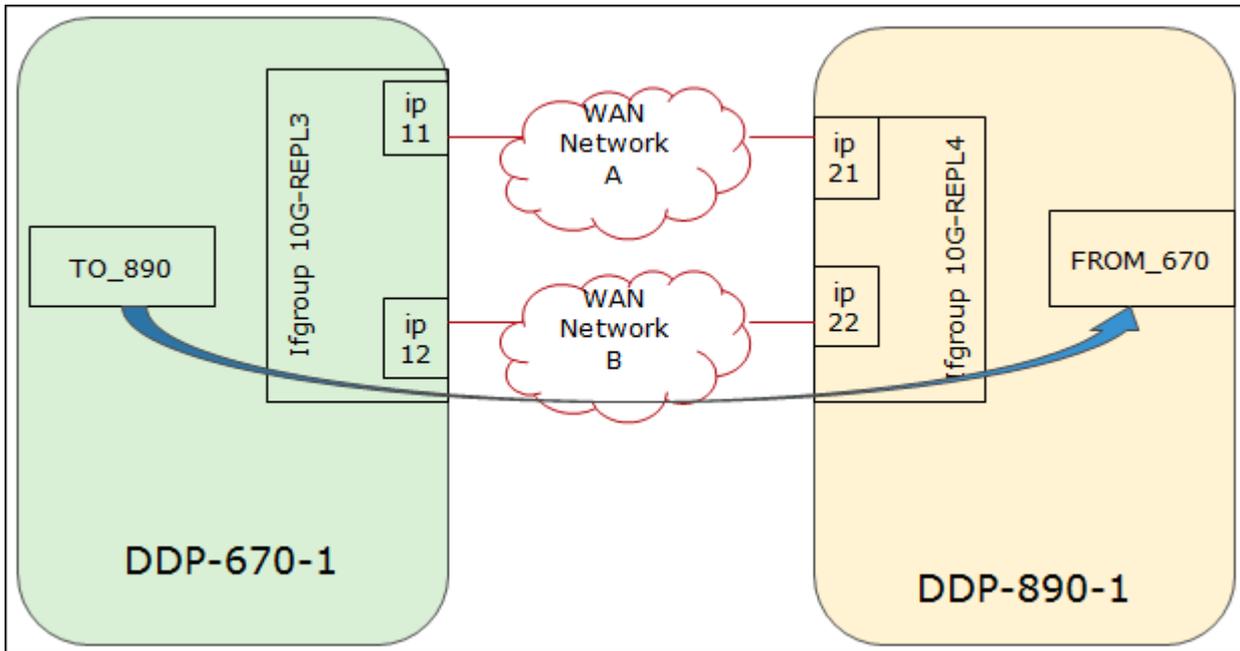
Replication over WANs

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 14](#) on page 55 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 14 Using Interface Groups for Replication over WANs

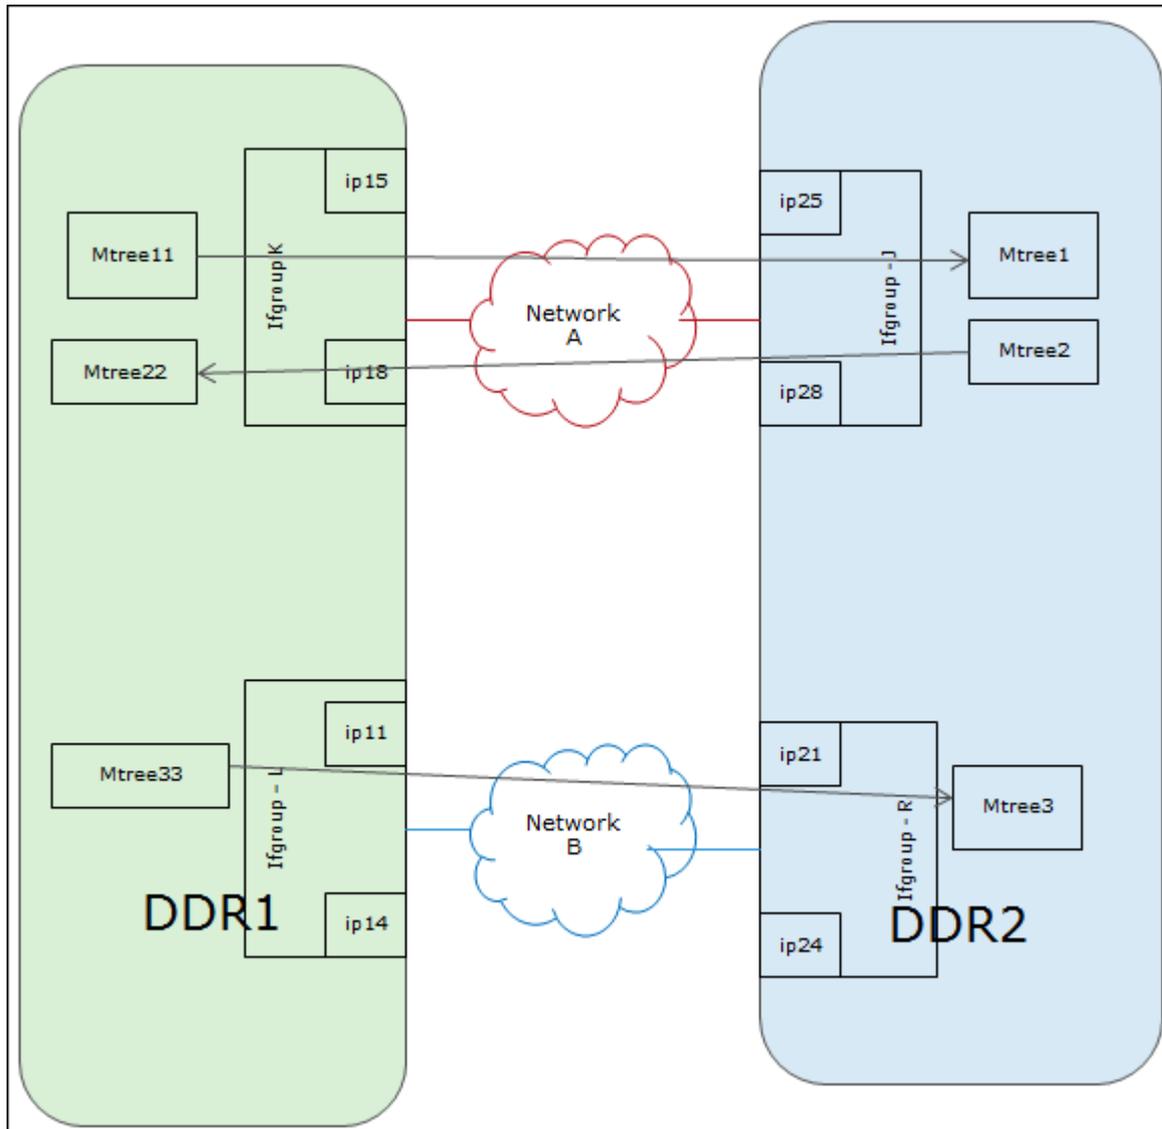
Other Supported Use Cases

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 15](#) on page 56, 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 15 Using Interface Groups for Replication over Multiple Customer Networks

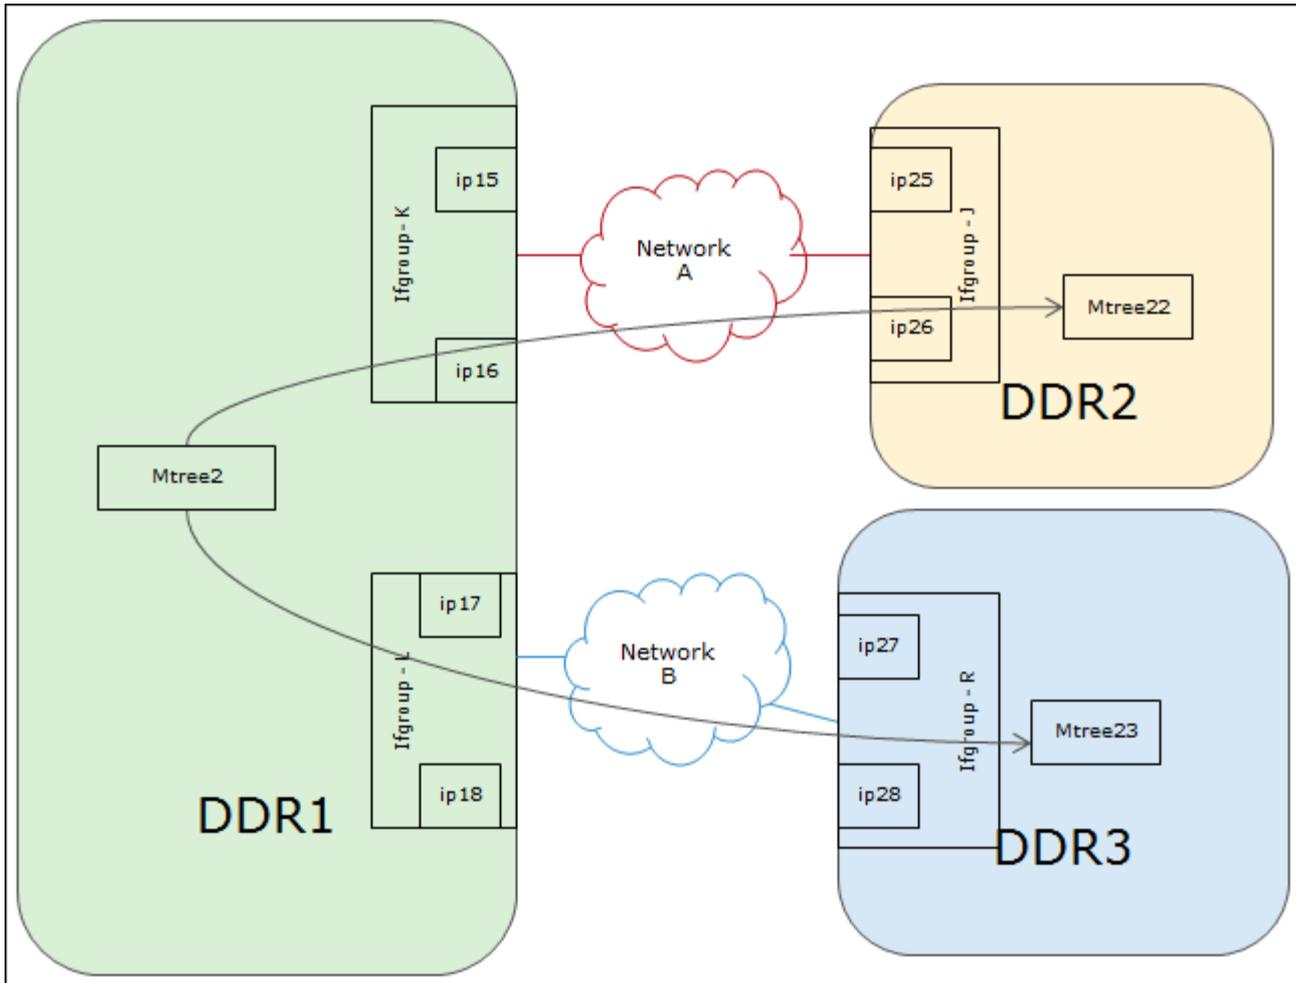


These replication paths have the same source and destination Data Domain systems, so the interface group is configured based on the MTree paths. For Network A, Ifgroup-K is configured based on the MTree11 and MTree22 paths, and Ifgroup-J is configured based on the MTree1 and MTree2 paths. For Network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

Fan-Out

In [Figure 16](#) on page 57, MTree2 is replicated from Data Domain System 1 (DDR1) to System 2 (DDR2), and from System 1 to System 3 (DDR3).

Figure 16 Using Interface Groups for Fan-Out Replication

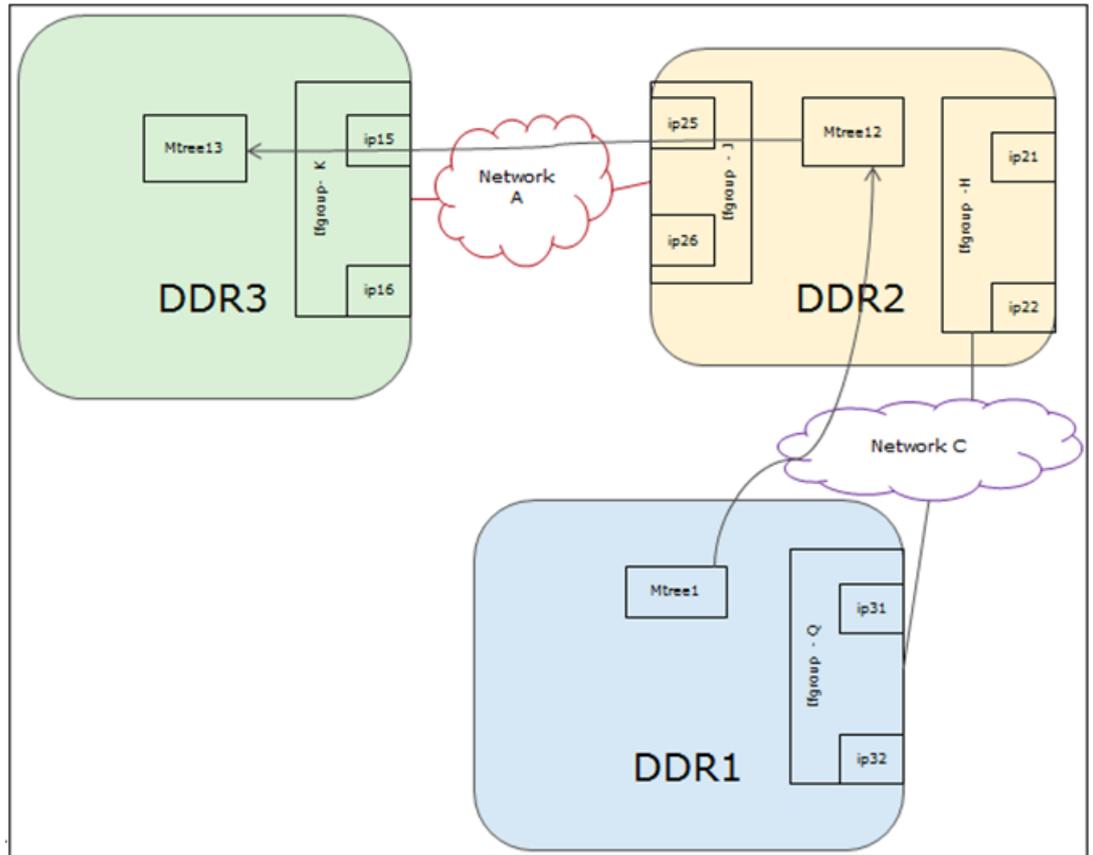


Because the same source MTree is used, the interface group is configured based on the hostname of the remote Data Domain system. For Network A, Ifgroup-K is configured based on the hostname of Data Domain System 2, and Ifgroup-J is configured based on the hostname of Data Domain System 1. For Network 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 17](#) on page 58, MTree1 is replicated from Data Domain System 1 (DDR1) to System 2 (DDR2), then to System 3 (DDR3).

Figure 17 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 28 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 18](#) on page 60, 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.

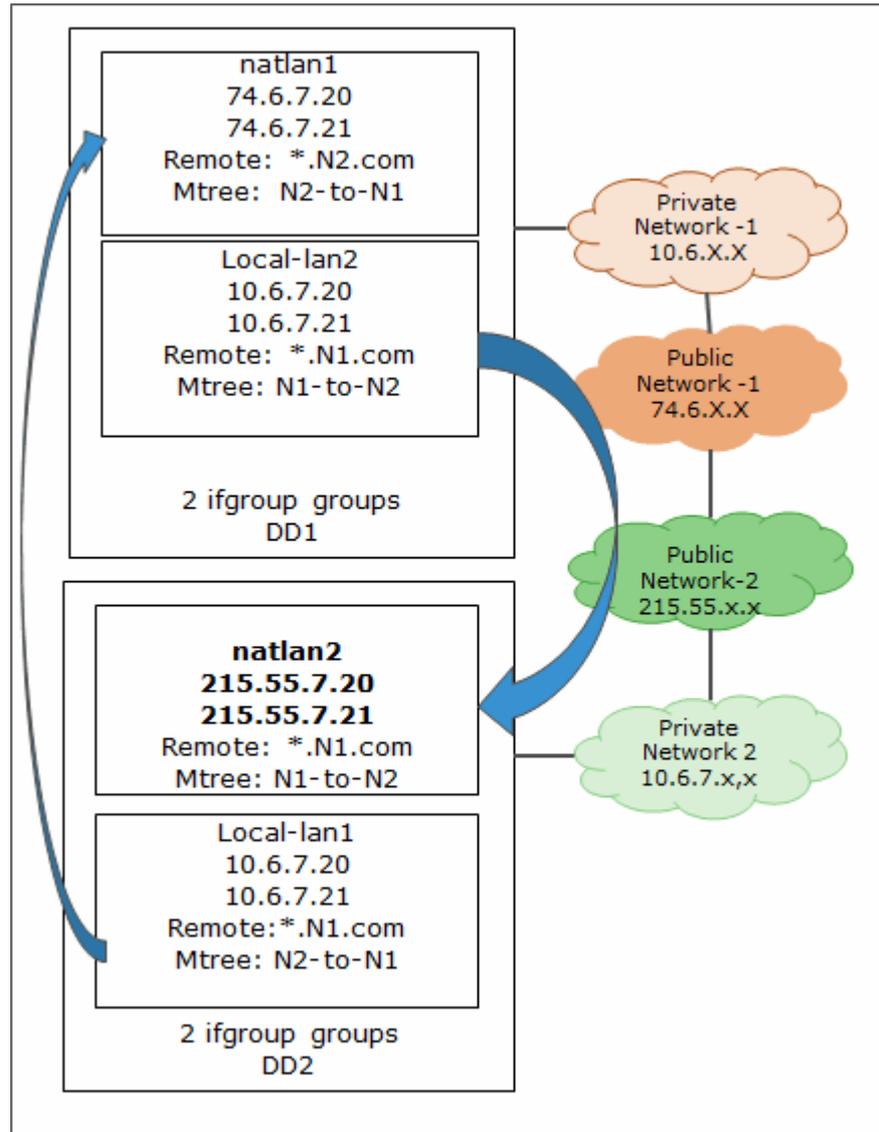
In [Figure 18](#) on page 60, DD1 and DD2 are each in a private network. Replication with NAT can be used in both directions, where the MTree name selects between the two interface groups. Only the source Data Domain system uses an interface group to select the source and destination IP. The source Data 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 18 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 enabled
Low 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 *Data Domain Operating System Administration Guide*.)

No configuration changes are necessary on the media 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 media server as this feature is transparent to the backup applications NetBackup and Backup Exec. 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 *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 media server client list may contain both fully qualified domain names or short names. The media 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.domain.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, *.*domain.com* followed by *.*com*
- Perfect match of sent client name, for example, ddbboost-dl.*domain.com*

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

Enabling In-flight Encryption

Run the following command to enable in-flight encryption for backup and restore operations over a LAN:

```
# ddbboost clients add <client-list> [encryption-strength {medium | high} authentication-mode {one-way | two-way | anonymous}]
```

This command can enable encryption for a single client or for a set of clients. The authentication-mode option is used to configure the minimum authentication requirement. A client attempting to connect using a weaker authentication setting will be blocked. Both one-way and two-way authentication require the client to be knowledgeable of certificates.

One-Way Authentication

The DD Boost client requests authentication from the Data Domain server, and the Data Domain server sends the appropriate certificate to the DD Boost client. The DD Boost client verifies the certificate. The communication channel between the DD Boost client and the Data Domain server is encrypted.

Two-Way Authentication

The DD Boost client requests authentication from the Data Domain server using the server certificate. The Data Domain server also requests authentication from the DD Boost client using the client certificate. After authentication through an SSL handshake, the communication channel between the DD Boost client and the Data Domain server is encrypted.

Anonymous Authentication

No certificates are exchanged, but information is exchanged. After the SSL handshake, the communication channel between the DD Boost client and the Data Domain server is encrypted.

Note

This option does not apply to DD Boost-over-Fibre Channel (FC). If both IP and FC are in use, encryption can be enabled on IP connections.

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 *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.
- To enable DD Boost-over-FC on clients running AIX, you can install the AIX DDdfc device driver. The SCSI disk device is available for AIX. For details, see [Installing the AIX DDdfc Device Driver \(Optional for AIX Clients\)](#) on page 69.

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
```

Example:

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

Note

You can use the `disk` option for the `ddbboost fc group add` command if you want to create a DFC group uses the client's native disk driver. This option is supported for AIX, Solaris, and Linux systems.

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 initiators to a `ddbboost-type scsitar` 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 `scsitar` 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 DDdfc 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 media server? The answer depends upon several factors:

1. Is the media 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 (for example, 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 media server and Data Domain system

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

For example, if:

- Media 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 media 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 media 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 media server. A best practice is to advertise the same number of DFC devices to all initiators on the same media 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 media servers using DD Boost 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 media 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 media/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 media 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 Media 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 media servers. No additional configuration is required. Linux media servers are not queue-

depth constrained, so many connections can share the same DFC generic SCSI device with no performance impact.

Windows Media 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 media 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 media 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 media 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 media 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 media 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 media server.

Assume:

- 2 physical paths between the media 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 media 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 media 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 Media 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 media 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 Media Servers

For the AIX DDdfc 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

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 Media 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.

A user who needs to use Solaris DFC disk as a non-root user must be assigned "sys_devices" privileges.

If a user needs to use Solaris DFC disk, you can assign "sys_devices" privileges as shown in the following example:

```
# # usermod -K defaultpriv=basic,proc_exec,sys_devices userid
```

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 also use the `ddboost fc group add` command with the optional `disk` characteristic to enable DD-Boost-over-FC. If you choose the latter option, no additional software needs to be installed. AIX will then treat associated DFC devices as native disk drives. The driver is packaged with your OST plug-in software.

For more information on the `ddboost fc group add` command, see the *Data Domain Operating System Command Reference Guide*.

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

Note

To maintain backwards compatibility, if you have both DDdfc devices and hard-disk devices on the AIX client that are connected to the same Data Domain system, the DD Boost plug-in selects DDdfc devices first for communications until no more remain before automatically switching to hard-disk devices.

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 15

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

Installing DD Boost for OpenStorage

Note

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

This chapter covers the following topics:

- [Installation Overview](#) 76
- [Installing OST Plug-In for NetBackup](#) 77
- [Installing OST Plug-In for Backup Exec](#) 80
- [Tuning Windows Media Servers for Performance](#) 81
- [Uninstalling the Windows Plug-in](#) 81

Installation Overview

The overall steps for installing Data Domain Boost are as follows:

Procedure

1. Obtain the license required to enable DD Boost on the Data Domain system.
 - The basic license allows you to back up and restore data.
 - A separate replication license enables you to perform Managed File Replication or Automatic Image Replication. You must obtain a replication license for both the source and destination Data Domain systems.
2. Enable and configure DD Boost on the Data Domain system. At a minimum, you must configure a DD Boost username and create a DD Boost storage unit.
 - To configure a DD Boost username, use the `ddboost user assign` command.
 - To create a DD Boost storage unit, use the `ddboost storage-unit create` command.
3. Install the OST plug-in software on each media server.

Note

- To enable DD Boost-over-FC on clients running AIX, you can install the AIX DDdfc device driver. For details, see [Installing the AIX DDdfc Device Driver \(Optional for AIX Clients\)](#) on page 69.
-

4. After you complete the installation steps described in this chapter, configure DD Boost as described in the chapter [Preparing the Data Domain System for DD Boost](#) on page 41.

OST Plug-In and DD OS Upgrades

Before upgrading either the OST plug-in or the DD OS, consult the *Data Domain Boost Compatibility Guide*. That guide specifies which versions of the OST plug-in are compatible with which versions of DD OS.

NOTICE

Failure to check the *Data Domain Boost Compatibility Guide* could result in unexpected incompatibilities requiring additional upgrades to correct.

To take advantage of new features in a DD OS release, upgrade the Data Domain system to the appropriate DD OS release, then upgrade the OST plug-in to a corresponding version. Although an older version of the OST plug-in maintains compatibility with a newer version of DD OS, it does not have the new functionality available in the newer version of the DD OS. Perform the upgrade as described in [Installing the OST Plug-In on Media Servers](#).

Note

This document illustrates the DD Boost configuration on Data Domain using commands in DD OS 5.7. If you are using a different version of DD OS with this version of the OST plug-in, see the corresponding DD OS Command Reference.

Firewalls and Ports

Note

This discussion applies only to DD Boost-over-IP.

The Data Domain system as it is initially configured does not work through a firewall (a media server to a Data Domain system, or from one Data Domain system to another). If you need the Data Domain system to work in the presence of a firewall, contact your network support provider.

The following ports must be open in a firewall for DD Boost backups and optimized duplication to work:

- TCP 2049 (NFS)
- TCP 2051 (required for Managed File Replication but not needed for Automatic Image Replication)
- TCP/UDP 111 (NFS portmapper)

Installing OST Plug-In for NetBackup

This section describes the commands used to install an OST plug-in within a NetBackup environment.

NetBackup environments consist of media servers and a master server. The master server manages clients and media servers and can also function as a media server. The OST plug-in must be installed on each media server. If a master server is also configured as a media server, then the OST plug-in must also be installed on the master/media server.

Note

Commands that run on the command line can be entered on either the master or the media server. If you run commands from the master server, use the `-media_server` option to tell NetBackup where to direct the operation that queries the plug-in about the properties of the storage server.

This guide uses the NetBackup commands located in the following directories, which you must add to your UNIX or Windows PATH.

Procedure

1. Add these directory locations to the UNIX PATH:

```
$ export
PATH=$PATH:/usr/opensv/netbackup/bin:
/usr/opensv/netbackup/bin/admincmd:/usr/opensv/volmgr/bin
```

2. Add these directory locations to the Windows PATH:

```
$ PATH=%PATH%;C:\Program Files\Veritas\NetBackup\bin;
C:\Program Files\Veritas\NetBackup\bin\admincmd;
C:\Program Files\Veritas\Volmgr\bin
```

Installing the OST Plug-In on Media Servers

The OST plug-in software must be installed on media servers that need to access the Data Domain system. When you upgrade the UNIX OST plug-in, the previous version of the plug-in is overwritten; therefore, you do not have to remove it. There are no special instructions to uninstall the OST plug-in on UNIX systems.

Installing the UNIX Plug-In

Procedure

1. Download the latest version of the OST plug-in from Online Support. Verify the MD5 digest of the download to assure its integrity by using an MD5sum or digest utility.
2. Enter `gunzip` or an equivalent command to unzip the tar file. Save the file on the media server in a location of your choice.
3. Stop the Remote Manager and Monitor Service (`nbrmms`) process of the backup application if it is running. Enter:

```
# nbrmms -terminate
```

4. Install the OST plug-in (a set of libraries in a `gtar` package.)
5. Use the `tar` command to uncompress the file:

```
# tar -vxf filename
```

6. The package also contains an installation script called `install.sh`, which verifies whether or not `nbrmms` has been stopped before you start the installation. Enter:

```
# install.sh -v verbose
```

If `-v` is not selected, information will only be written to `/log/OST/EMC/logs/install-datestring.log`.

The shared library files that the script installs are `libstspiDataDomain.so` and `libstspiDataDomainMT.so`.

If the plug-in automatically exists, it is automatically overwritten.

A script called `uninstall.sh` is now available if you want to remove existing plug-ins.

7. If the plug-in already exists, you are prompted to enter `y` to proceed.
8. Restart the backup application's `nbrmms` process by entering:

```
# nbrmms
```

Installing the Windows Plug-In

The Windows plug-in installer is `libstspiDataDomainSetup.exe`. It supports 64-bit Windows plug-ins.

Preparing for Installation

Procedure

1. Download the latest version of the Windows OST plug-in installer from Online Support.

2. Verify the MD5 digest of the download to assure its integrity, by using WinMD5 or similar utility. Unzip the plugin to extract `libstspiDataDomainSetup.exe`.
3. Stop any NetBackup services. Follow the instructions given in [Starting, Stopping, and Restarting NetBackup Windows Services](#) to stop the service.
4. Remove any previous plug-in version by using the Windows Control Panel or by executing the OST plug-in uninstall command in silent mode.

```
libstspiDataDomainUninstall.exe /S
```

or interactive mode:

```
libstspiDataDomainUninstall.exe
```

Starting the Installation

You can run the installation in an interactive mode or in silent mode.

Procedure

1. Double-click the set-up executable to launch the installer.
The installer determines whether NetBackup is installed and whether its respective services are running. If the installer detects a service that is running, it displays a message to this effect and exits.
2. If the services have been stopped, the installer displays the license agreement. Read the terms and click **I Agree** to continue.
3. In the Choose Install Location dialog box, the correct destination folder is shown. Do not change this folder. Click **Install** to start the installation.
A progress bar monitors the installation.
4. When the Installation is complete, you can click the **Show details** button to view the files installed and their location.
5. Restart all services. See [Starting, Stopping, and Restarting NetBackup Windows Services](#) on page 79.
6. Tune the Windows media server for performance. See [Tuning Windows Media Servers for Performance](#) on page 81.

NetBackup Services

Follow the instructions for starting, stopping, and restarting UNIX or Windows services.

Starting and Stopping NetBackup UNIX Services

To stop UNIX services, enter:

```
# nbrmms -terminate
```

To start or restart UNIX services, enter:

```
# nbrmms
```

Starting, Stopping, and Restarting NetBackup Windows Services

Procedure

1. Go to **Start > Control Panel > Administrative Tools > Services**.

2. In the Services window, services are listed in alphabetical order by name. Locate the name **NetBackup Remote Manager and Monitor Service**. Its `Status` field shows the state of the service.
3. Select the service and right-click.
4. The menu that displays has options to **Stop**, **Start**, or **Restart** the service. Select the appropriate menu option.

Installing OST Plug-In for Backup Exec

Installing the Plug-In on Media Servers

The OST plug-in software must be installed on media servers that need to access the Data Domain system. Because Backup Exec supports OpenStorage only on Windows media servers, the following section covers instructions for Windows servers only.

Note

Backup Exec is not supported with DD Boost-over-FC.

Install the Windows Plug-In

The Windows plug-in installer is `libstspiDataDomainSetup.exe`. This supports 64-bit Windows plug-ins.

Procedure

1. Prepare for installation.
 - a. Download the latest version of the Windows OST plug-in installer from Online Support.
 - b. Stop any Backup Exec services. Follow the instructions given in [Backup Exec Services](#) to stop the service.
 - c. Remove any previous plug-in version using either the Windows Control Panel or `libstspiDataDomainUninstall.exe`.
 - d. Double-click the set-up executable to launch the installer. The installer determines whether Backup Exec is installed and whether its respective services are running. If the installer detects that a service is running, it displays a message to this effect and exits.
 - e. Proceed to start the installation.
2. Start the installation.
 - a. If the services have been stopped, the installer displays the license agreement. Read the terms. Select **I Agree** to continue.
 - b. In the Choose Install Location dialog box, the correct destination folder is shown. Do not change the folder. Select **Install** to start the installation.

Note

A progress bar monitors the installation.

- c. When the Installation is complete, you can select the **Show details** button to view the files installed and the location of those files.

- d. Restart all Backup Exec services. See [Backup Exec Services](#).
3. Verify that the Backup Exec Deduplication Option is enabled.

Backup Exec Services

Starting, Stopping, or Restarting Windows Services

Within Backup Exec, you can start, stop, and restart Backup Exec Services in the Backup Exec Service Manager window.

Upgrading the Backup Exec Plug-In

Procedure

1. Run the Backup Exec Services Manager and stop all services (do not close the dialog box).
2. Remove the old plug-in and install the new plug-in.
3. Return to the Backup Exec Services Manager dialog box to restart all services.
4. Close the dialog box.

Tuning Windows Media Servers for Performance

For tuning information, refer to the Knowledge Base article, *Tuning Windows Media Servers for Performance*, Document ID 000180974, which is available at [Online Support](#).

Uninstalling the Windows Plug-in

This procedure applies to NetBackup and Backup Exec.

Procedure

1. Stop the services of the backup application.
2. Do one of the following:
 - Uninstall the Data Domain OpenStorage plug-in in Windows' Control Panel uninstall/remove program feature (as you would uninstall a typical Windows program).
 - Double-click `libstspiDataDomainUninstall.exe`, which was installed in the same directory as the plug-in. Click **Uninstall**. After the uninstall, click **Show details** to view which files were removed.

CHAPTER 5

Backup Application Administration

Note

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

This chapter covers the following major topics:

- [Configuring a Media Server](#) 84
- [NetBackup Administration](#) 98
- [Backup Exec Administration](#) 100

Configuring a Media Server

Media server configuration depends on the backup application being used. See the appropriate configuration section.

NetBackup Configuration

Note

The examples in this chapter assume the following configuration:

- A media server with the name `load64` that runs NBU 7.6
 - Two Data Domain systems with DD Boost enabled named `dd22` and `dd100`.
-

Media server configuration consists of the following procedures:

- Registering each Data Domain system
 - Scan for newly added devices, especially for Boost-over-FC
 - Adding credentials for each media server that is to communicate with a Data Domain system
 - Creating disk pools
 - Creating storage units, which are collections of disk pools
 - Setting backup policies
-

Note

Commands that run on the command line can be entered on either the master or the media server. If you run commands from the master server, use the `-media_server` option to tell NetBackup where to direct the operation that queries the plug-in about the storage server's properties.

Concurrent Connection Limit

With Backup Exec, the maximum number of concurrent connections (jobs) from a single media server is 64.

Registering Each Data Domain System

Procedure

1. On the media server, start the backup application's services. See [NetBackup Services](#).
2. On the media server, verify that the plug-in is detected by the backup application by entering:

```
# bpstsinfo -pi -stype DataDomain
```

The output shows:

- the vendor version, which is the plug-in version
 - the build version, which is the OST plug-in version.
3. On the Data Domain system, enable virtual synthetics if that feature is planned to be used, by entering the following command:

```
# ddbboost option set virtual-synthetics enabled
```

4. On the media server, register a Data Domain system by entering:

For DD Boost-over-IP:

```
# nbdevconfig -creatests -stype DataDomain -storage_server dd22 -media_server load64
```

In this case, the `storage_server` can be either an IP address or a hostname, such as `dd22`.

For DD Boost-over-FC:

```
# nbdevconfig -creatests -stype DataDomain -storage_server DFC-dd100 -media_server load64
```

In this case, the `storage_server` prefix `DFC-` indicates the desire to use the DDBoost-over-FC transport to communicate with the Data Domain system. The name following the prefix is the DFC-server-name of the desired Data Domain system, such as `dd100`.

5. Repeat the above procedure for each Data Domain system that will be running DD Boost.

Adding Credentials

Procedure

1. On a media server that needs to communicate with a Data Domain system, enter:

```
# tpconfig -add -storage_server dd22 -stype DataDomain -sts_user_id username -password password
```

Note

NetBackup versions 7.6 and later allow the credentials to also be configured from within NetBackup. See the NetBackup documentation for more information.

Note

The `ddbboost storage-unit create storage-unit user user-name` command is now available for each storage-unit to be distinct from one another.

2. Repeat the above step for each media server that needs to communicate with a specific Data Domain system. The following is an example for DFC server using `dd100`:

```
# tpconfig -add -storage_server DFC-dd100 -stype DataDomain -sts_user_id username -password password
```

Results

After you add the credentials, the backup application does the following:

- Saves the credentials so the media server can log into the Data Domain system.
- Configures the media server as a data mover that can transfer data between the primary storage (the backup application's client) and the storage server (the Data Domain system). The backup application maintains an access path between the media server and the storage server.

Creating Disk Pools

Disk pools are collections of disk volumes that the backup application administers as single entities. Disk pools correspond to storage units.

Note

Each disk pool requires a unique name.

The backup application provides a command line interface (CLI) and a graphical user interface (GUI). You can use either to create disk pools.

Creating a Disk Pool

Procedure

1. The backup application's Remote Manager and Monitor Service (`nbrmms`) must be running. To start it, enter:

```
# nbrmms
```

2. Obtain the identity of the storage unit on the Data Domain system (`dd22`) by entering:

```
# nbdevconfig -previewdv -storage_server dd22-stype DataDomain
> /tmp/dvlist
```

3. Create a disk pool using the information obtained from the previous command by entering:

```
# nbdevconfig -createdp -dp dd22_storage-unit1_dp -stype
DataDomain
-storage_servers dd22 -dvlist /tmp/dvlist
```

The disk pool name must be unique.

Output similar to the following is displayed:

```
Disk pool dd22_storage-unit1_dp has been successfully created
with 1 volume.
```

Creating Storage Units

A storage unit contains a disk pool. Multiple storage units can be grouped together into a Storage Unit Group. You can create storage units using either the CLI or the GUI.

Note

Each storage unit requires a unique name.

Creating Storage Units

Procedure

1. Enter a command similar to the following:

```
# bpstuadd -label dd22_storage-unit1_su -dp dd22_storage-
unit1_dp
-host load64a -M load64a
```

Note

There is no output from this command.

Creating a Backup Policy

For instructions on creating a backup policy, see the NetBackup 7.x Administration Guides.

Configuring Buffers

You can set the number and size of various buffers, but you cannot change their size limits. The location for these files depends on your operating system.

- The UNIX file location is `/usr/opensv/netbackup`.
- The Windows file location is `install_path\netbackup\db\config`.

For best performance, set `SIZE_DATA_BUFFERS` and `SIZE_DATA_BUFFERS_DISK` to 262144.

To set the number and size of buffers, create the following files, as appropriate for your operating system.

- `NET_BUFFER_SZ`
 - Description: TCP/IP socket buffer size
 - Media: N/A
 - Default on UNIX: 32,032
 - Default on Windows: 32,032

- `NUMBER_DATA_BUFFERS`

Note

The number must be a power of two.

- Description: Number of shared data buffers.
 - Media: Tape
 - Default on UNIX: 8/4 (Non-multiplexed/multiplexed.)
 - Default on Windows: 16/8 (Non-multiplexed/multiplexed.)
- `NUMBER_DATA_BUFFERS_RESTORE`
 - Description: Number of shared data buffers.
 - Media: Tape
 - Default on UNIX: 8/12 (Non-multiplexed/multiplexed.)
 - Default on Windows: 16/12 (Non-multiplexed/multiplexed.)

- `NUMBER_DATA_BUFFERS_DISK`

Note

The number must be a power of two.

- Description: Number of shared data buffers.
 - Media: Disk
 - Default on UNIX: 8/4 (Non-multiplexed/multiplexed.)
 - Default on Windows: 16/8 (Non-multiplexed/multiplexed.)
- `SIZE_DATA_BUFFERS`

Note

The size must be a multiple of 32 KB. The default used when this file does not exist is 32 KB. The maximum value supported by the Data Domain plug-in is 1 MB. The default value when the file exists, and the recommended value for best performance is 256 KB.

- Description: Size of shared data buffers.
 - Media: Tape
 - Default on UNIX: 64 KB
 - Default on Windows: 64 KB
 - `SIZE_DATA_BUFFERS_DISK`
-

Note

The size must be a multiple of 32 KB. The default used when this file does not exist is 32 KB. The maximum value supported by the Data Domain plug-in is 1 MB. The default value when the file exists, and the recommended value for best performance is 256 KB.

- Description: Size of shared data buffers.
- Media: Disk
- Default on UNIX: 256 KB
- Default on Windows: 256 KB
- `SIZE_DATA_BUFFERS_NDMP`
 - Description: Buffer size for NDMP backups.
 - Media: N/A
 - Default on UNIX: 63 KB
 - Default on Windows: 63 KB

Configuring Optimized Duplication

The OST plug-in enables a NetBackup media server to specify a duplication process and delegate its execution to the Data Domain system. This sharing has the following advantages:

- The backup application system retains control of creating and duplicating backup files and keeps track of all copies in its catalog, which ensures easy and efficient recovery.
- Optimized duplication removes the media server from the data path in creating duplicate copies of backup images, which reduces the load on the backup application system and frees it for other work.
- The Data Domain system uses Wide Area Network (WAN) efficient replication process for deduplicated data. The process is optimized for WANs, which reduces the overall load on the WAN bandwidth required for creating a duplicate copy.
- Data Domain Replicator software features, such as the Low-Bandwidth Optimization Option, can be used transparent to the backup application to reduce further the data sent over WAN links that are fewer than 6 Mb/s.
- Data Domain Replicator software features, such as Encrypted Optimized Duplication, are transparent to the backup applications. This feature allows all data

that is sent over the WAN for the purpose of creating duplicate copies to be encrypted, which provides higher security.

It is recommended that you add the destination Data Domain system's IP address to the source Data Domain system using the `net hosts add ipaddr {host | "alias host"}... command`.

Note

All media servers, source and destination, must have permission to access both Data Domain systems. You should add all of the media servers that need to access a Data Domain system to it using the `net hosts add command`.

DD Boost-Over-Fibre Channel Considerations

DD Boost-over-FC introduces a complication to the procedure for configuring optimized duplication.

An optimized duplication operation requires communication between three systems:

- Media_Server
- Src_DD_System — The source Data Domain system
- Dst_DD_System — The destination Data Domain system

During an optimized duplication operation, the Dst_DD_System is accessed by both of the other systems:

- By Media_Server — for control operation/setup
- By Src_DD_System — for data transfer

The Media_Server-to-Dst_DD_System communication may use either of the following transports:

- DD Boost-over-IP
- DD Boost-over-FC

But the Src_DD_System-to-Dst_DD_System communication is always via IP networking.

Now, consider the case where the Media_Server uses DD Boost-over-FC to communicate with the Dst_DD_System. The full optimized duplication operation now requires two "names" for the Dst_DD_System:

- DFC-*<dfc-server-name>* -- needed by DD Boost Library on the Media_Server
- IP hostname -- needed by the Src_DD_System

However, during configuration, only a single name for Dst_DD_System is presented to the DD Boost Library: the DFC-style name, DFC-*<dfc-server-name>*.

The DD Boost Library has to pass a name to the Src_DD_System as part of the request to start transferring the data.

The Src_DD_System needs an IP hostname for the Dst_DD_System, since all communication between the two Data Domain systems is performed using IP networking.

But the DD Boost Library knows the Dst_DD_System only by its DFC-style name. So, what name for the Dst_DD_System should the DD Boost Library present to the Src_DD_System?

The answer is that the DD Boost Library just strips off the "DFC-" prefix, and presents the Dst_DD_System's DFC-server-name to the Src_DD_System.

For example:

```
Media Server: clientA
Src_DD_System: DFC-ddr1
Dst_DD_System: DFC-ddr2
```

In this case, the DD Boost Library will present to the Src_DD_System the name `ddr2` as the Dst_DD_System.

This works naturally if Dst_DD_System's DFC-server-name is the same as its IP hostname, as known to Src_DD_System. This is the expected normal situation, since the default DFC-server-name for a Data Domain system is its simple nodename.

If the user has changed Dst_DD_System's DFC-server-name to something else (e.g., `my-ddr-via-dfc`), then he needs to make sure that when Src_DD_System performs a hostname lookup of `my-ddr-via-dfc`, it finds an IP address by which Dst_DD_System is reachable. This can be achieved by adding an entry to the `/etc/hosts` file on Src_DD_System.

Using Storage Lifecycle Policies to Automate Optimized Duplication

A storage lifecycle policy consists of a list of destinations for backup files and a retention period for each file. A lifecycle process creates, retains, and finally expires the files. Using storage lifecycle policies allows you to specify different retention periods for the initial backup and for the duplicate copies. For example, you might specify one retention period for the original local backup and another for a duplicate at a disaster recovery site.

Select individual storage unit as duplication destination in SLP. For further information, refer to the Knowledge Base article, OST Duplication Does Not Work, Document ID 000181560, which is available on the Online Support site <https://support.emc.com>.

Note

- If there is a preferred link or IP address for sending the optimized duplication data between two Data Domain storage servers, use that link or address when creating the destination storage server.
 - Should you ever want to start optimized duplication manually, use the NBU CLI command `bpduplicate`, which is described in the Veritas NetBackup documentation.
-

Configuring a Virtual Synthetic Backup

To use virtual synthetic backups, set up the policy attributes and schedules as follows:

Procedure

1. In DD OS 5.4, Virtual Synthetics is enabled by default. If it is disabled, enable a virtual synthetic backup on the Data Domain system by entering:

```
# ddboost option set virtual-synthetics enable
```

2. Verify that NetBackup has enabled virtual synthetics on the Data Domain system and verify that the `OptimizedImage` flag is set by entering:

```
# nbdevquery -liststs -U
```

Note

If you are using an old disk pool created before DD OS 5.2 using DD Boost Version 2.5, then `ddboost option set virtual-synthetics enable` command will not work as intended. The job will finish but you will not find the above messages as NetBackup does regular synthetic replication. In such a case, perform the following steps:

- a. Create a new disk pool in NetBackup.
- b. Add the flag manually to the existing disk pool, by entering the following command:

```
# disk-pool-name: dlh35-dp
storage-server-name: dlh35

nbdevconfig -changests -storage_server dlh35 -stype DataDomain -
setattribute OptimizedImage

nbdevconfig -changedp -dp dlh35-dp -stype DataDomain -setattribute
OptimizedImage
```

- c. Verify that the flag `OptimizedImage` is added to the disk pool using the following command:

```
# nbdevquery -listdp -U -dp dlh35-dp
```

If the `OptimizedImage` flag is not displayed in the output, configure it with the `nbdevconfig` command:

```
# nbdevconfig -changests
```

Sample Backup Operations

The following examples show the commands to initiate backups and display various types of backups.

Sample Backup Operation: Full Backup

A full backup will consist of a header (HDR) image file, one or more fragment (F1) image files and a true image restore (TIR) image file as can be seen on the DDR storage unit.

```
# ddbboost storage-unit show sparcl compression

List of files in sparcl and their compression info:

rtp-ost-sparcl.datadomain.com_1309959523_C1_HDR:
1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 8.9
  Original Bytes:          8,924
  Globally Compressed:     8,924
  Locally Compressed:      767
  Meta-data:               236

rtp-ost-
sparcl.datadomain.com_1309959523_C1_F1:1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 1.0
  Original Bytes:          931,228,244
  Globally Compressed:     927,741,488
  Locally Compressed:     942,139,003
  Meta-data:               3,091,380

rtp-ost-sparcl.datadomain.com_1309959523_C1_TIR:
1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 43.9
  Original Bytes:          100,349
  Globally Compressed:     54,304
  Locally Compressed:      1,912
  Meta-data:               376
```

Sample Backup Operation: Incremental Backup

An Incremental backup will add a header (HDR) image file, one or more fragment (F1) image files and a true image restore (TIR) image file as can be seen on the DDR storage unit as shown in bold below.

```
# ddbost storage-unit show sparcl compression

List of files in sparcl and their compression info:

rtp-ost-sparcl.datadomain.com_1309959523_C1_HDR:
1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 8.9
    Original Bytes:          8,924
    Globally Compressed:     8,924
    Locally Compressed:      767
    Meta-data:               236
rtp-ost-
sparcl.datadomain.com_1309959523_C1_F1:1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 1.0
    Original Bytes:          931,228,244
    Globally Compressed:     927,741,488
    Locally Compressed:      942,139,003
    Meta-data:               3,091,380
rtp-ost-sparcl.datadomain.com_1309959523_C1_TIR:
1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 43.9
    Original Bytes:          100,349
    Globally Compressed:     54,304
    Locally Compressed:      1,912
    Meta-data:               376
rtp-ost-sparcl.datadomain.com_1309959822_C1_HDR:
1309959822:dd670c2-1:4:0:::
Total files: 1; bytes/storage_used: 8.8
    Original Bytes:          8,924
    Globally Compressed:     8,924
    Locally Compressed:      776
    Meta-data:               236
rtp-ost-
sparcl.datadomain.com_1309959822_C1_F1:1309959822:dd670c2-1:4:0:::
Total files: 1; bytes/storage_used: 93.9
    Original Bytes:          931,227,936
    Globally Compressed:     9,784,959
    Locally Compressed:      9,890,654
    Meta-data:               28,684
rtp-ost-sparcl.datadomain.com_1309959822_C1_TIR:
1309959822:dd670c2-1:4:0:::
Total files: 1; bytes/storage_used: 39.3
    Original Bytes:          100,528
    Globally Compressed:     66,592
    Locally Compressed:      2,151
    Meta-data:               404
```

Sample Backup Operation: Synthetic Full Backup

The synthetic full will add a header (HDR) image file, one or more fragment (F1) image files and a true image restore (TIR) image file as can be seen on the DDR storage unit as shown in bold below.

```
# ddbost storage-unit show sparcl compression

List of files in sparcl and their compression info:

rtp-ost-sparcl.datadomain.com_1309959523_C1_HDR:
1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 8.9
    Original Bytes:          8,924
    Globally Compressed:     8,924
    Locally Compressed:      767
    Meta-data:               236
rtp-ost-
```

```

sparc1.datadomain.com_1309959523_C1_F1:1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 1.0
  Original Bytes:          931,228,244
  Globally Compressed:    927,741,488
  Locally Compressed:     942,139,003
  Meta-data:              3,091,380
rtp-ost-sparc1.datadomain.com_1309959523_C1_TIR:
1309959523:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 43.9
  Original Bytes:          100,349
  Globally Compressed:    54,304
  Locally Compressed:     1,912
  Meta-data:              376
rtp-ost-sparc1.datadomain.com_1309959822_C1_HDR:
1309959822:dd670c2-1:4:0:::
Total files: 1; bytes/storage_used: 8.8
  Original Bytes:          8,924
  Globally Compressed:    8,924
  Locally Compressed:     776
  Meta-data:              236
rtp-ost-
sparc1.datadomain.com_1309959822_C1_F1:1309959822:dd670c2-1:4:0:::
Total files: 1; bytes/storage_used: 93.9
  Original Bytes:          931,227,936
  Globally Compressed:    9,784,959
  Locally Compressed:     9,890,654
  Meta-data:              28,684
rtp-ost-sparc1.datadomain.com_1309959822_C1_TIR:
1309959822:dd670c2-1:4:0:::
Total files: 1; bytes/storage_used: 39.3
  Original Bytes:          100,528
  Globally Compressed:    66,592
  Locally Compressed:     2,151
  Meta-data:              404
rtp-ost-sparc1.datadomain.com_1309959823_C1_HDR:
1309959823:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 8.9
  Original Bytes:          8,924
  Globally Compressed:    8,924
  Locally Compressed:     768
  Meta-data:              236
rtp-ost-
sparc1.datadomain.com_1309959823_C1_F1:1309959823:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 1.0
  Original Bytes:          7,435,452
  Globally Compressed:    7,420,935
  Locally Compressed:     7,444,262
  Meta-data:              23,812
rtp-ost-sparc1.datadomain.com_1309959823_C1_TIR:
1309959823:dd670c2-1:4:1:::
Total files: 1; bytes/storage_used: 43.0
  Original Bytes:          100,449
  Globally Compressed:    54,304
  Locally Compressed:     1,958
  Meta-data:              376

```

The synthetic backup is done using the DDP_SYNWR API which can be displayed on the Data Domain system by the `ddbost show stats` and `ddbost show histograms` commands.

```

# ddbost show stats
07/06 07:13:38

DD Boost statistics:

...
DDP_SYNWR          :          18          [0]
...

Count              Errors

```

Image creates	9	0
Image deletes	0	0
Pre-compressed bytes received	3,712,802,816	-
Bytes after filtering	1,856,586,752	-
Bytes after local compression	1,856,586,752	-
Network bytes received	1,857,697,928	-
Compression ratio	2.0	-
Total bytes read	0	0

Configuring an Auto Image Replication Backup in a Source Domain

Backups must be directed to a disk pool that has a Source Replication properly configured.

A Storage Lifecycle Policy (for example, AIR-test1-test2) is created. It contains a standard Backup step and a Duplication (NBU 7.1) step specifying “Remote master (send to the replication target device in a remote domain).” A policy is created specifying the SLP as its policy storage.

Configuring an Auto Image Replication Backup in a Target Domain

Backups to be automatically imported must be file-copied to a disk pool that has a Destination Replication property.

A Storage Lifecycle Policy (for example, AIR-test1-test2) named identically to the one in the Source domain is created. It contains an import step.

When Targeted AIR is used, the source and the destination master media servers' Storage Lifecycle Policies must have the “data classifications” set to the same color code (for example, silver, gold, or platinum).

Running an Auto Image Replication Backup

When an Auto Image Replication backup runs in the Source domain, there is a backup step followed in time (by default, 30 minutes later) by a duplication step.

AIR replication job count will be displayed as a `Src-repl` job in the output of a `ddboost show connections` command, the same as other NetBackup and Backup Exec optimized duplication jobs.

After the duplication in the Source domain, some time later (again by default, 30 minutes), the imported image-set is available as shown in the Activity Monitor of the Target domain.

Unlike other NetBackup and Backup Exec optimized duplication jobs, AIR replication jobs will not be displayed as a `Dst-repl` job in the output of a `ddboost show connections` command.

Backup Exec Configuration

Note

DD Boost-over-Fibre Channel is not supported with Backup Exec.

For information on setting up, scheduling, and monitoring jobs, see the *Veritas Backup Exec 15 Administrator's Guide*.

For all Backup Exec versions, complete the following steps:

Procedure

1. Create a logon account with the following information.

- a. Non-Default Login account.
- b. DD Boost user-name
- c. DD Boost password

Note

The `ddbost storage-unit create storage-unit user user-name` command is now available for each storage-unit to be distinct from one another.

2. Configure devices.
 - a. Create a storage unit on the Data Domain system.
 - b. Add an OpenStorage server specifying the Data Domain host name and the logon account name previously created.
 - c. Backup Exec will query the Data Domain system for a list of storage-units. Select a storage unit.
 - d. Specify the number of concurrent operations for the device. The total number of concurrent connections (jobs) from a single media server OpenStorage plug-in to all associated OpenStorage storage units is 48. The concurrent operations limit for a single device can be determined as follows: $48 \geq \# \text{ OpenStorage storage units} + \sum \text{ concurrent operations for each storage unit}$ In the case of a single Data Domain system with a single storage unit, the concurrent operation count can be set as high as 47.
 - e. Specify the default values for Disk Space Management. The data stream chunk size ranges from 64 KB to 256 KB. For best performance, 256 KB is recommended.
 - f. Specify the Storage unit sharing value. A single Data Domain storage unit can be shared by multiple media servers when the shared media servers are associated with a single primary media server. In the media servers list, select the primary media server.
 - g. Restart the Backup Exec services when a new Data Domain system is added.

Create a Logon Account

Follow these steps to create a logon account.

Procedure

1. Double-click the icon to the left of **1) Create Logon Accounts** in the Getting Started panel of the Home page. The Logon Account Wizard Welcome dialog box is displayed. Click **Next**.
2. In the Set Up a Logon Account dialog box, select **Add a new logon account**, and click **Next**.
3. In the Enter Logon Account Credentials dialog box, enter the user name and password set for DD Boost. Click **Next**.
4. In the Logon Account Name dialog box, type an account name that describes this logon account. Click **Next**.
5. In the Type of Logon Account dialog box, make the account available to all Backup Exec users. Click **Next**.
6. In the Default Logon Account dialog box, select **No**. The Data Domain system account is usually not the Backup Exec system logon. Click **Next**.

7. Verify your account settings as shown in the Logon Account Summary dialog box. Click **Back** to edit prior selections. If the account information is correct, click **Next**.
8. The Completing the Logon Account Wizard dialog box is displayed. Click **Finish**.

Configuring Devices

Follow these steps to configure devices.

Procedure

1. Create a storage unit on the Data Domain system. See [Creating Storage Units](#).
2. From the Backup Exec Home page, select **Storage > Configure Storage > Network Storage > Next > OpenStorage**.

Note

Backup Exec's Deduplication option must be installed and licensed for the OpenStorage option to be available.

3. Configure the Add OpenStorage Device dialog box's General tab as follows:
 - **Name:** Enter the name of the Data Domain system.
 - **Server:** Enter the Data Domain host name.
 - Select the logon account name previously created.
 - Select **DataDomain** as the server type.
 - **Storage unit:** Select storage unit.
 - **Concurrent Operations:** Specify the number of concurrent operations for the device.
 The total number of concurrent connections (jobs) from a single media server OST plug-in to all associated OpenStorage storage units is 48. The concurrent operations limit for a single device can be determined as follows:

$$48 \geq \# \text{ OpenStorage storage units} + \sum \text{ concurrent operations for each storage unit}$$
 In the case of a single Data Domain system with a single storage unit, the concurrent operation count can be set as high as 47.
4. Click **OK**.
5. Configure the Add OpenStorage Device dialog box's Advanced tab as follows:
 - Accept the default values for **Disk Space Management** and **Direct Access**.
 - Specify a **Data stream chunk size** from 64 KB to 256 KB. For best performance, 256 KB is recommended.
6. Click **OK**.
7. Click the **Sharing** tab.

A single Data Domain storage unit can be shared by multiple media servers when the shared media servers are associated with a single primary media server.

In the media servers list, select the primary media server, and click **OK**.

8. You must restart the Backup Exec services when a new Data Domain system is added. In the Restart Services dialog box, click **Restart Now**.

Results

After the device has been configured, the new storage unit is displayed in the Devices page.

Configuring Optimized Duplication

The ways to develop duplication jobs in Backup Exec are described in detail in the *Veritas Backup Exec 15 Administrator's Guide*. You can attach an associated duplicate job to any backup job, or duplicate a previous backup set.

The OST plug-in enables a media server to specify a duplication process and delegate its execution to the Data Domain system. This sharing has the following advantages:

- The backup application system retains control of creating and duplicating backup files and keeps track of all copies in its catalog, which ensures easy and efficient recovery.
- Optimized duplication removes the media server from having to create duplicates of backup files, which reduces the load on the backup application system and frees it for other work.
- The Data Domain system uses Wide Area Network (WAN) efficient replication process for deduplicated data. The process is optimized for WANs, which reduces the overall load on the WAN bandwidth required for creating a duplicate copy.
- Data Domain Replicator software features, such as Low-Bandwidth Optimization Option, can be utilized transparent to the backup application for further reducing the data sent over WAN links that are less than 6 Mb/s.
- Data Domain Replicator software features, such as Encrypted Optimized Duplication, can be used transparent to the backup applications. This feature allows all data sent over the WAN for the purpose of creating duplicate copies to be encrypted, which provides higher security.

A best practice is to add the destination Data Domain system's IP address to the source Data Domain system using the following command:

```
net hosts add ipaddr {host | "alias host"}
```

Note

All media servers, source and destination, must have permission to access both Data Domain systems. It is recommended that you add all of the media servers that need to access a Data Domain system to it using the `net hosts add` command. To duplicate an image from one system to another, the following conditions must be met:

- The Data stream chunk size for devices configured on both Data Domain systems between which optimized duplication is to take place must be set to the same value. It is recommended that this value be 256 KB as shown in the OpenStorage Device Properties dialog box.
- The Concurrent Operations count of the destination Data Domain system is greater than or equal to that of the source Data Domain system.

Configuration Limitations for Optimized Duplication

- Optimized Duplication is supported with Backup Exec 2010 R2 or higher.
- Data Domain supports optimized duplication for images that have only one dataset. If multiple volumes or selections from multiple volumes (C:\Windows,

D:, E:, etc.), or agents (SQL Server, SharePoint, etc.), or a combination are being backed up in one job, then the resulting backup image contains datasets for all the drives or the applications unless Veritas Backup Exec Hotfix 138226 is applied. This Hotfix can be applied only to Backup Exec 2010 R2. With Hotfix 138226 applied, Backup Exec creates multiple images, one for each dataset in the backup job. In the above example that contains multiple volumes in a job, there would be three images produced—one for C:\Windows, one for D: and one for E:. Optimized duplication of select individual images, or all three images, can then be carried out by Backup Exec.

NetBackup Administration

Find your OST Plug-in Version

Procedure

1. Enter:

```
# bpstsinfo -pi -stype DataDomain
```

Results

The output shows the vendor version, the plug-in version, and the build version.

Find your NetBackup version

Procedure

1. Display by entering:

```
# cat <NetbackupInstall_Dir>/version
```

Results

Sample output:

```
Netbackup-Solaris10 7.6
```

Network Time-Outs

Backup and restore jobs often take a long time to complete. Although the OST plug-in 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.

A best practice is 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.

Set Backup Application Time-out Using the CLI

Procedure

1. Add the following two lines to the `<NetBackupInstall_directory>/bp.conf` file:

```
CLIENT_CONNECT_TIMEOUT = 1800
CLIENT_READ_TIMEOUT = 1800
```

Note

The time-out value is expressed in seconds.

Set Backup Application Time-out Using the GUI

Procedure

1. Expand the **NetBackup Management** node.
2. Expand **Host Properties**.
3. Select **Master Servers**.
4. In the right pane, double-click the machine name.

In the property dialog box that is displayed, change the time-out values.

Grouping Storage Units to Provide Failover

The administrator can specify a group of storage units to share a workload. The administrator tells the backup application system how to choose among the storage units in the group for the next job by setting one of the following selection criteria:

- Failover (This is the recommended setting)
Setting failover as the selection criterion ensures that a backup job does not fail if the storage unit to which it is directed fails. The backup application chooses another storage unit in the same group to finish the job.
- Prioritized
- Round robin
- Load balance

Delete a Data Domain storage server

NOTICE

This procedure removes all of the data and resources associated with the storage server. Do not attempt this procedure unless it is necessary.

Procedure

1. Delete all of the files specified by the `BACKUP_ID` by entering:

```
# bpexpdate -backupid BACKUP_ID -d 0
```

2. Delete all of the policies from the GUI.

3. Delete all of the storage units by entering:

```
# bpstudel -label SU_NAME
```

4. Delete all the disk pools by entering:

```
# nbdevconfig -deletedep -stype DataDomain -dp pool-name
```

5. Delete the storage server by entering:

```
# nbdevconfig -deletests -storage_server dd22 -stype DataDomain
```

Note

You can use the GUI to delete the files, lifecycle policies, storage units, and disk pools.

For troubleshooting information, see [Unable to Delete the Data Domain System](#).

6. Remove the credential using the `tpconfig` command.

```
# tpconfig -delete -storage_server dd22 -stype DataDomain -  
sts_user_id username
```

Backup Exec Administration

Find your OST plug-in version

Procedure

1. Go to the Backup Exec install directory and find the file `libstspiDataDomain.dll`.
2. Right-click the file's name and select **Properties** from the menu.
3. Select the **Details** tab. The OST plug-in version is displayed as the file version.

Find your Backup Exec version

Procedure

1. From the Backup Exec Home page, select **About** from the Help menu.

Delete Storage Units on Data Domain Systems

Procedure

1. There are two options for deleting a storage unit on a Data Domain system:
 - You can erase all media within a Backup Exec device (a Data Domain system's storage unit) and then delete the device from Backup Exec.
 - You can also delete the device from Backup Exec even if media remains in the device. The storage unit remains on the Data Domain system and some files are left in the storage unit. To recover this space, delete the storage unit on the Data Domain system by entering:

```
# ddbost storage-unit delete storage-unit
```

CHAPTER 6

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 [Online Support](#).

This chapter covers the following topics:

- [General Troubleshooting](#).....102
- [Data Domain System Settings for File Replication](#)..... 102
- [NetBackup Troubleshooting](#)..... 102
- [Backup Exec Troubleshooting](#)..... 113

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 *Data Domain Boost Compatibility Guide* at <http://compatibilityguide.emc.com:8080/CompGuideApp/>.

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. [Configuring a Media Server](#) provides instructions on establishing user 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 optimized duplications to fail for NetBackup and Backup Exec.

NetBackup Troubleshooting

Unable to Delete the Data Domain System

This procedure assumes the following:

- You are unable to delete the Data Domain system.
- You have already run the `nbdevconfig` command with the `deletests` option and it has failed, which means that the `emm` or `rmms` process might be down.
- All of the files for the specified Data Domain have expired. For instructions on how to expire a file, see your NBU documentation.

If you are still unable to delete the Data Domain system, follow these steps:

Procedure

1. Enter:

```
# nbdevconfig -deletests -storage_server DDR -stype DataDomain
```

2. If core files result, contact Data Domain Support. Otherwise, continue to the next step.
3. Follow the instructions below for your operating system.

On a Windows System

Procedure

1. Restart the NetBackup services on the media server by running these two executable files:

```
NBUInstallPath\NetBackup\bin\bpdown.exe
NBUInstallPath\NetBackup\bin\bpup.exe
```

2. Run `deletests` again. If it fails, enable more detailed NBU logging by opening the `NBUInstallPath\NetBackup\nblog.conf` file and adding this entry:

```
NBSTSI=OID=202
```

3. Enable detailed logging messages on media servers as described in [Error Logging on the Media Servers](#).

On a UNIX System

Procedure

1. If `rmms` restarts but `emm` does not, verify that all of the processes are up, especially `emm` or `rmms`.

2. If these processes are not up, enter:

```
# /bp/bin/goodies/netbackup start
```

3. Run `deletests` again. If it still fails, enable more NBU logging by opening the `/bp/nblog.conf` file and adding this entry:

```
NBSTSI=OID=202
```

4. Enable detailed logging messages as described in [Error Logging on the Media Servers](#).

Check the Installation

Problems with basic operations such as backups may result from improper installation.

Procedure

1. Verify that the files are in the correct location by entering the following, depending on your operating system:

- a. On a UNIX system, enter:

```
# ls /usr/openv/lib/ost-plugins/
```

The command output should include the names of the shared library files:

```
libstspiDataDomain.so
libstspiDataDomainMT.so
```

Note

You can also check the install log found at `/log/OST/EMC/logs/install-datestring.log`.

- b. On a Windows system, enter:

```
C:\Program Files\Veritas\bin\ost-plugins
```

The command output should be the name of the shared library file `libstspiDataDomain.dll`.

- Determine the plug-in version by entering:

```
# bpstsinfo -pi
```

The vendor version shown in the output is the Data Domain plug-in version, and build version is the version of the DD Boost API.

Note

If the `bpstsinfo` command fails, check the log files in the `/usr/opensv/netbackup/logs/admin` directory.

Check Credentials

Procedure

- To display credentials for all Data Domain systems registered as storage servers, enter the following command from the backup application system:

```
# tpconfig -dsh -all_hosts -stype DataDomain
```

After you finish

If you receive a message stating that you failed to add credentials for the Data Domain system (OpenStorage server), follow the procedure [Adding Credentials](#), which describes how to set up credentials and check for errors and inconsistencies.

Resolve License Errors

If the Configure Disk Pool wizard reports a license error, do the following:

Procedure

- Open the file `bp.conf`.
- Check if it contains an extra `CLIENT_NAME` entry.
- Delete any extra `CLIENT_NAME` entry.

Error Logging on the Media Servers

The error log is the main tool for troubleshooting problems related to NetBackup in an OpenStorage environment.

Procedure

- Before starting a backup, restore, or optimized duplication operation, enable logging on the NetBackup media server. Follow the instructions for the media server's operating system, or use the NetBackup GUI.

- Enable error logging on a UNIX system

Enter:

```
# /usr/opensv/netbackup/logs/mklogdir
```

- Enable error logging on a Windows system

Enter:

```
C:\Program Files\Netbackup\logs\mklogdir.bat
```

Results

After you have enabled logging, the OST plug-in prefixes error and informational log messages with the name `DataDomain`.

Resolving Failed Backups on Media Servers

Search for plug-in error messages in the log file as described below for the media server's operating system.

Resolve Failed Backups on a UNIX System

Procedure

1. Enter:

```
# cat /usr/opensv/netbackup/logs/bptm/LOGFILE_DATE | grep
DataDomain
```

The command selects lines from the specified log file that contain the word `DataDomain`. The plug-in uses `DataDomain` as a prefix for its log messages.

Resolve Failed Backups on a Windows System

Procedure

1. Enter:

```
C:\Program Files\Veritas\logs\bptm\LOGFILE_DATE.log
```

2. Open the log file and search for the word `DataDomain`.

Resolve Failed File Duplication

Procedure

1. Search for plug-in error messages in the media server log files, which are specific to the server's operating system:

- UNIX

- For `read_file`:

```
/usr/opensv/netbackup/logs/bpdm
```

- For `write_file`:

```
/usr/opensv/netbackup/logs/bptm
```

- For `file-replication`:

```
/usr/opensv/netbackup/logs/bpdm
```

- Windows

- For `read_file`:

```
C:\Program Files\Veritas\logs\bpdm
```

- For `write_file`:

```
C:\Program Files\Veritas\logs\bptm
```

- For `write_file`:

```
C:\Program Files\Veritas\logs\bptm
```

2. Verify that the replication license is installed by entering:

```
# license show
```

3. For further assistance, contact your contracted support provider.

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:

```
# fileysys status
```

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

```
# nfs status
```

Resolve Plug-In Log Messages

When the plug-in encounters an error, it returns an `EPLUGIN` error code to NetBackup and logs a reason for the error.

Procedure

1. Determine if the reason is one of the following:
 - `Write Length Exceeds Limit Error`
The write buffer data size is limited. If you receive an exceeds limit error message, change the buffer size to a value within the specified limit as described in [Configuring Buffers](#).
 - `Program Not Registered`
The following output indicates that the program is not registered:

```
(: RPC: Program not registered)
```

2. Enable DD Boost by installing a valid license:

```
# license add ddboost-license-code
```

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

```
# fileysys status
```

Resolve “Cannot connect on socket” Error

This error results when the command `nbdevconfig -creatests` has been run, but the storage server is not created because of a socket connection error.

Follow these steps:

Procedure

1. Check to make sure the `nbemm` process is running. If it keeps failing upon startup, usually there is an issue with the NBU database.
2. Use the `vxlogview` utility to check the logs located in `/usr/opensv/logs/51216-*.log` for errors.
3. Recreate the Database. Enter:

```
# /usr/opensv/db/bin/create_nbdb -drop
```

NetBackup Backup Jobs Fail on Solaris Media Servers

If a file backup job fails with a media write error (84) at the start of the job, a typical activity monitor job detail might contain the following:

```
2/28/2009 3:36:22 AM - Critical bptm(pid=1750) failure to open sts
for storage server apodddrrp01: plug-in reports error 2060046 plugin
error2/28/2009 3:36:23 AM - end writing media open error(83)
```

The `bptm` log may contain information similar to the following:

```
01:33:02.585 [28874] <16> apoddrp01: /usr/opensv/lib/ost-plugins/
libstspiDataDomain.so:stspi_open_server STS_EPLUGIN Can't connect to
mountd on apoddrp01 (: RPC: Miscellaneous tli error - An event
requires attentionError 0)
```

In the above example, an entry in `/etc/inet/ipsecinit.conf` has enforced encryption on traffic from port 665 (`sun-dr`). However, the Solaris operating system had Sun Dynamic reconfiguration disabled. As a result, although the media server used port 665 to connect via NFS to the Data Domain system, the packet did not leave the media server because it was not encrypted.

To fix this problem, you need to disable dynamic reconfiguration.

Disable dynamic reconfiguration

Procedure

1. Uncomment or remove `sun-dr` entries in `/etc/inet/inetd.conf`:

```
sun-dr stream tcp wait root /usr/lib/dcs dcssun-dr stream tcp6
wait root /usr/lib/dcs dcs
```

2. Have `inetd` reread the configuration file, by entering:

```
kill -HUP pid-inetd
```

3. Uncomment or remove the `sun-dr` entries in `/etc/inet/ipsecinit.conf`:

```
{dport sun-dr ulp tcp} permit {auth_algs md5}{sport sun-dr ulp
tcp} apply {auth_algs md5 sa unique}
```

4. Remove the active IPsec configuration from the running system.

- a. Obtain the index numbers by entering:

```
ipseccconf | grep sun-dr
```

- b. Delete the policy for `sun-dr` by entering:

```
ipseccconf -d index
```

Backup jobs from AIX clients fail

If you run multiple backup jobs from an AIX client to an HA-enabled Data Domain system with DD Boost, the backup jobs might fail to finish.

Traditionally, DD Boost allocates 24 MB of client-side buffers for each open file. However, when backups are sent to an HA-enabled DD system, clients require 128 MB for each open file. This requirement can create problems on any client that opens a large number of files but is low on memory.

AIX clients can be particularly susceptible to this issue due the size of the default heap in AIX applications.

You can obtain better results by setting the AIX environment variable to `export LDR_CNTRL=MAXDATA=0x70000000`.

Optimized Duplication Job Fails

The replicator software license for optimized duplication is required on both the source and destination Data Domain systems that run DD OS 4.7 or later.

If this license is not installed, an optimized duplication job fails. A typical activity monitor job detail indicates a media write error (84) occurred. The NetBackup `bpdm` log states that the NFS operation is not supported.

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
```

Virtual Synthetic Backup

- Verify that normal backups are OK.
- Verify that the Storage Lifecycle Policy attributes are set properly.
- Verify that TIR files are being generated in the storage unit.

```
# ddboost storage-unit show [compression] [storage-unit] [tenant-unit tenant-unit]
```

- Verify that DDP_SynWR RPCs are being sent.

```
# ddboost show stats
```

- Verify that OptimizedImage flag is set.

```
# nbdevquery -liststs
```

- Verify virtual-synthetics is enabled on the Data Domain system.

```
# ddboost option show
```

Monitoring Auto Image Replication

On the source Data Domain system, statistics and histograms are reported for RPCs directly related with Auto Image Replication: DDP_REMFILEOPS and DDP_IMAGESETOPS. Also DDP_IMAGESETS is a count of all image-sets sent from this Data Domain system. The DDP_IMAGESETS histogram reports the time from the last image of the image-set being sent for file-copy until the event is posted on the Data Domain system in the target domain.

On the target Data Domain system, statistics and histograms are reported for the DDP_GETEVENT RPC. Also DDP_EVENTS is a count of all image-set events reported for import. The DDP_EVENTS histogram reports the time from which the event was posted on the Data Domain system in the target domain until it is delivered to NetBackup for import.

Use the `ddboost file-replication show` commands to get the file-replication performance reports of individual files.

To display the DD Boost statistics, enter:

```
# ddboost show stats
```

To display the DD Boost histogram, enter:

```
# ddboost show histogram
```

Auto Image Replication Not working

Procedure

1. To verify that the connection from the source Data Domain system to the target Data Domain system for replication is working, enter:

```
# replication option show
```

Note

Make sure the TCP port is 2051 or as set.

2. To verify that the associations are properly configured at the source and target Data Domain systems, enter:

```
# ddbboost association show
```

3. To verify new backup images on the source Data Domain system, enter:

```
# ddbboost storage-unit show source-su
```

4. To verify new backup images on the target Data Domain system, enter:

```
# ddbboost storage-unit show target-su
```

Note

Make sure that the image names on the target are identical to those on the source.

- a. If new file-copied images are not being sent, check for file-copy errors and problems reported on the source Data Domain system.
 - b. Elevate the debug level for the `bpdm log` and inspect it for problems.
5. To verify the statistics on the target Data Domain system, enter:

```
ddbboost show stats
```

Note

If the target Data Domain system shows that `DDP_GETEVENT` total count is increasing with no errors in the target domain, this indicates that some NetBackup target domain on this Data Domain system is periodically polling for events.

Note

If `DDP_GETEVENT` is not increasing, enter:

```
nbdevconfig -updatestats -storage_server  
rtp-ost-dd670c2.datadomain.com -stype DataDomain
```

Restart NetBackup services.

6. Confirm that the disk pools used the report properly. If they do not, update the database, enter:

```
nbdevconfig -updatestats -storage_server  
rtp-ost-dd670c2.datadomain.com -stype DataDomain
```

Restart NetBackup services.

- a. On NBU 7.1, look for `STS_LSUF_REP_TARGET` and `_SOURCE` flags, enter:

```
bpstsinfo -li -stype DataDomain -sn rtp-ost-  
dd670c2.datadomain.com
```

LSU Info:

```
Server Name: DataDomain:rtp-ost-dd670c2.datadomain.com  
LSU Name: sparcl2sol02  
Allocation : STS_LSU_AT_STATIC  
Storage: STS_LSU_ST_NONE  
Description: Data Domain SU for DDBOOST images  
Configuration:  
Media: (STS_LSUF_DISK | STS_LSUF_ACTIVE |  
STS_LSUF_REP_ENABLED | STS_LSUF_REP_TARGET)
```

```

Save As : (STS_SA_OPAQUEF)
Replication Sources: 1
( ddp-890-1.datadomain.com:sparc12sol02 )
Replication Targets: 0 ( )
Maximum Transfer: 1048576
Block Size: 32768
Allocation Size: 0
Size: 8295733002240
Physical Size: 8295733002240
Bytes Used: 40263745536
Physical Bytes Used: 40263745536
Resident Images: 0
LSU Info:
Server Name: DataDomain:rtp-ost-dd670c2.datadomain.com
LSU Name: sol022sparc1
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: Data Domain SU for DDBOOST images
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_ACTIVE |
STS_LSUF_REP_ENABLED | STS_LSUF_REP_SOURCE)
Save As : (STS_SA_OPAQUEF)
Replication Sources: 0 ( )
Replication Targets: 1
( ddp-890-1.datadomain.com:sol022sparc1 )
Maximum Transfer: 1048576
Block Size: 32768
Allocation Size: 0
Size: 8295733002240
Physical Size: 8295733002240
Bytes Used: 40263745536
Physical Bytes Used: 40263745536
Resident Images: 0

```

b. On NBU 7.5 and 7.6, look for Replication Target and Source, enter:

```

nbdevquery -listdp -stype DataDomain -U
Disk Pool Name      : sol022sparc1-dd670c2
Disk Pool Id       : sol022sparc1-dd670c2
Disk Type          : DataDomain
Status             : UP
Flag               : Patchwork
Flag               : Visible
Flag               : OpenStorage
Flag               : SingleStorageServer
Flag               : CopyExtents
Flag               : AdminUp
Flag               : InternalUp
Flag               : LifeCycle
Flag               : CapacityMgmt
Flag               : FragmentImages
Flag               : Cpr
Flag               : FT-Transfer
Flag               : OptimizedImage
Flag               : ReplicationSource
Raw Size (GB)      : 7726.00
Usable Size (GB)   : 7726.00
Num Volumes        : 1
High Watermark     : 98
Low Watermark      : 80
Max IO Streams     : -1
Comment            :
Storage Server     : ost-dd670c2.datadomain.com (UP)

```

7. To check that .imgset and event files are on the target Data Domain system, enter:

```
ddboost event show target-su
```

Note

The `.imgset` files are named with a form:

```
192:rtp-ost-sparc1.datadomain.com_ddr1.domain1.com_1328637954_1.imgset
```

Where: `192:rtp-ost-sparc1.datadomain.com` is the NetBackup image set name consisting of the job number and source client host name (with any embedded `_` converted to `-`). `ddr1.domain1.com` is the hostname of the source Data Domain system. `1328637954` is the NetBackup image timestamp (in this case the image was created 2/7/12 18:05. For Excel, the timestamp in A2 is converted to a time/date by the formula `=A2/86400+DATE(1970,1,1)` 1 is the number of images in the set - currently always `1.imgset` is the identifier.

Note

The `.event` files are named with a form:

```
bluemia.datadomain.com_31234_6589_1.event.0000000000000006
```

Where: `bluemia.datadomain.com` is the hostname of the NetBackup media server that first detected the associated imageset. `31234` is the process ID of the NetBackup media server. `6589` is the thread ID of the NetBackup media server. `1.event 0000000000000006` is the unique event identifier on this Data Domain system.

- a. The long term presence (more than two hours) of `.imgset` files in the target storage unit indicates that the target NetBackup domain is not querying for posted events. DD OS removes any `.imgset` files that are more than two hours old.
 - b. The long term presence (more than two hours) of event files in the target storage unit indicates that the target NetBackup domain is not processing events. This may mean that the SLP specified in the `.imgset` file is not spelled correctly in the target NetBackup domain.
8. Confirm that the NetBackup database reflects that the plug-in is an event source and that the `DDP_GETEVENT` RPC count using the `ddboost show stats` command is incrementing. If not, update the database using

```
nbdevconfig -updatestats -storage_server
rtp-ost-d670c2.datadomain.com -stype DataDomain
```

- a. Look for `STS_SRV_EVSYNC`, enter:

```
bpstsinfo -si -stype DataDomain -sn rtp-ost-
dd670c2.datadomain.com
```

Server Info:

```
Server Name: DataDomain:rtp-ost-dd670c2.datadomain.com
Supported Stream Formats:
[
]
Server Flags: (STS_SRV_IMAGELIST | STS_SRV_CRED |
STS_SRV_EVSYNC | STS_SRV_IMAGE_COPY)
Maximum Connections: 149
Current Connections: 0
Supported Interfaces:
[
]
Supported Credentials:
[
]
```

b. Elevate the unified log debug level for stsem to 6

```
vxlogcfg -a -p 51216 -o stsem -s DebugLevel=6
```

c. Capture a time period and review the stsem log for errors specifying a start date and time:

```
vxlogview -p 51216 -o stsem -b "2/7/2012 3:30:00 PM" > c:\stsem.log
```

Or capture previous number of hours:

```
vxlogview -p 51216 -o stsem -t 4 > c:\stsem.log
```

d. Return the unified log debug level for stsem to 1 so that the logs do not fill the file system:

```
vxlogcfg -a -p 51216 -o stsem -s DebugLevel=1
```

9. In the stsem log, look for a log entry indicating that the event has been posted for import. Once posted for import the event file is deleted from the target Data Domain system.

```
02/22/12 07:05:17.307 [STSEventSupplier::postReplEvent()]
AddOstImageToImport seqno=52 masterServer=
<rtp-ost-sparc1.datadomain.com> media=<rtp-ost-
sparc1.datadomain.com>
origin_NBU_master=<bluemedia> isi_slpname=<AIR-vol1-vol2>
e_orig_server=<DataDomain:rtp-ost-dd670c2.datadomain.com>
e_num_images=<1> : [0] servername=<rtp-ost-
dd670c2.datadomain.com>
servertype=<DataDomain> imo_lsu.sln_name=<vol2>
imo_def.img_basename=<bluemedia_1329923106_C1_IM>
```

10. If ddbboost event show indicates that events are being consumed at the target domain (no events listed for the given target storage unit) but the activity monitor does not show Import activity, verify that the times on the source and target domain media master servers are reasonably synchronized (typically within a few minutes or less). They do not have to be in same time zone.
11. In the bpcd log of the destination Data Domain media server performing the AIR import, look for an entry indicating the import job has been started. This can most easily be done by grepping for the image ID reported at the source domain. In this case, bluemedia_1329923106.

```
07:05:29.624 [24145] <2> process_requests: fork cmd =
/usr/opensv/netbackup/bin/bpdm bpdm -restoretir -cmd -b
bluemedia_1329923106 -c
bluemedia -cn 1 -drn 0 -p @aaaaal -v -jobid 285 -from_replica -
mst 6
```

12. In the bpdbm log, the following log entries are found.

```
07:05:33.463 [24169] <2> db_logimagerec: backup id
bluemedia_1329923106
```

13. Finally in the bpdm log, the import takes place:

```
7:05:30.232 [24150] <2> bpdm: INITIATING (VERBOSE =
5): -restoretir -cmd -b bluemedia_1329923106 -c
bluemedia -cn 1 -drn 0 -p @aaaaal -v -jobid 285 -from_replica -
mst 6
```

14. If the import job is failing with no images were successfully processed (191) message, please review the detail information in the bpimport log. In the display below the SLP on the target domain did not match the SLP in the source domain.

```
04/06/2012 11:23:38 - Error bpimport (pid=11137)
Import of replica image, backup id ostqa-
```

```
sparc1.datadomain.com_1333724024,
Import failed because the imported image specifies an SLP name
which does not exist
```

15. Detailed logging of AIR operations on the Data Domain system is available in `ddfs.info` if the proper level of debugging is enabled. When running at a default (level 0) debugging level, `ddfs.info` contains error messages from `ost_remfileops`, `ost_imagesetops`, `ost_get_event_id`, `ost_get_event`. These indicate catastrophic errors. In order to see more typical errors, the debug level (-D) of the OST debug mask (-M) needs to be elevated to 3. This can be done using the `config_debug` utility:

```
/ddr/bin/config_debug -P /ddr/bin/ddfs -M 0x00100000 -D 3
```

Note

Complete operation logging is available at debug level 6. However, debug level 6 is typically not used due to the volume of logging output. If debug level 6 is used, the debug level must be returned to 0 in `ddfs.info` after capturing the problem.

Cancel Auto Image Replication

To stop replications in progress for a given SLP, as suggested in the Veritas Best Practices Guide, enter:

```
nbstlutil cancel -lifecycle SLP name -force
```

Note

See the Veritas NetBackup Auto Image Replication FAQ at https://www.veritas.com/support/en_US/article.000112974.

Backup Exec Troubleshooting

Basic Troubleshooting

- Verify that the concurrent connections (jobs) count is set properly for all storage units.
 - Backup Exec: The total number of concurrent connections from a single media server plug-in to all associated OpenStorage storage units is 48. This number was specified when you configured the device. See [Configuring Devices](#) on page 96.
- When encountering a problem, try to stop Backup Exec services and restart them. If this does not work:
 - Reboot the server.
 - Start the debugger and try to recreate the problem.

Check the installation

Problems with basic operations such as backups may result from improper installation.

Procedure

1. Verify that `libstspiDataDomain.dll` is in `C:\Program Files\Symantec\Backup Exec\`.

Note

Although NetBackup is now a Veritas product, the use of Symantec in the file path found in the previous command is still correct.

2. Determine the plug-in version by right-clicking on the DLL and opening its **Properties > Details**.

Check Credentials for a Data Domain System

Procedure

1. Display the OpenStorage device properties noting the log on account.
2. Verify that the logon username matches the DD Boost username on the Data Domain system.

Resolve License Errors

Backup Exec needs to be licensed for OpenStorage which is part of the deduplication license option.

Set Up Active Debugging

Use the Backup Exec debugging utility (SGMON) to troubleshoot Backup Exec issues.

Procedure

1. Run the Backup Exec Debug Monitor for Active Debugging.
2. The following Capture options must be selected (enabled): **Job Engine**, **Backup Exec Server**, and **Device and Media**.
3. **Capture to file** must be enabled.
4. Set **Device and Media Debug** and select **Enable verbose logging**, if it is not enabled.