



# EMC<sup>®</sup> Data Domain Boost

Version 2.6

## Administration Guide

302-001-618

REV 01

Copyright © 2015 EMC Corporation. All rights reserved. Published in USA.

Published January, 2015

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

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

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

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

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

# CONTENTS

	<b>Preface</b>	<b>5</b>
<b>Chapter 1</b>	<b>Introducing Data Domain Boost</b>	<b>9</b>
	Revision History.....	10
	Overview of DD Boost.....	10
	Supported Configurations .....	11
	Upgrade Compatibility.....	11
<b>Chapter 2</b>	<b>DD Boost Features</b>	<b>13</b>
	Overview of DD Boost Features.....	14
	Distributed Segment Processing .....	14
	Managed File Replication.....	15
	Low-Bandwidth Optimization.....	15
	Encrypted Managed File Replication.....	15
	IFGROUP: DD Boost IP Load Balancing and Failover.....	16
	Interfaces.....	18
	Clients.....	18
	DD Boost-over-Fibre Channel Transport.....	19
	DD Boost-over-Fibre Channel Path Management.....	21
	Initial Path Selection.....	22
	Dynamic Re-Balancing.....	22
	Client Path Failover.....	23
	Queue-Depth Constraints.....	23
	Virtual Synthetic Backups.....	24
	Backup Server Access Validation.....	24
	Storage Unit Capacity Quotas.....	24
<b>Chapter 3</b>	<b>Preparing the Data Domain System for DD Boost</b>	<b>27</b>
	Enabling DD Boost on a Data Domain System.....	28
	Creating Storage Units.....	28
	Configuring Quotas for Storage Units (Optional) .....	28
	Configuring Distributed Segment Processing.....	29
	Configuring IFGROUP .....	29
	Modifying an Interface Group.....	31
	Removing an Interface Group.....	32
	Configuring Managed File Replication (MFR).....	33
	Throttling MFR.....	33
	Enabling Low-Bandwidth Optimization .....	33
	Enabling Encryption.....	33
	Changing the MFR TCP Port.....	34
	Configuring Media Server Access Validation.....	34
	Configuring DD Boost-over-FC Service.....	35
	Sizing DD Boost-over FC Device-Set.....	36
	Sizing Calculation.....	37
<b>Chapter 4</b>	<b>Administration</b>	<b>39</b>

	Configuring a Backup Server.....	40
	Backup Administration.....	40
	Network Time-Outs.....	40
<b>Chapter 5</b>	<b>Basic Troubleshooting</b>	<b>41</b>
	General Troubleshooting.....	42
	Resolve time-out error.....	42
	Managed File Replication Job Fails.....	42
	Add license for Replication.....	42
	Verify Encrypted Managed File Replication Configuration.....	42
	Virtual Synthetic Backup.....	43

# Preface

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

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

---

## Note

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

---

## Purpose

This guide explains how to install, configure, and use the EMC Data Domain Boost with Data Domain systems.

---

## Note

Always check the EMC Online Support site <https://support.emc.com> for the latest version of this document before proceeding. Any printed or CD-based version of this document may be out-of-date.

---

## Audience

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

## Related EMC documentation

The following Data Domain system documents provide additional information:

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

## DD Boost Backup Application Documentation

Documentation for backup applications is available through backup application vendor.

## Special notice conventions used in this document

EMC uses the following conventions for special notices:



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

---

**⚠ WARNING**

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

---

**⚠ CAUTION**

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

---

**NOTICE**

Addresses practices not related to personal injury.

---

**Note**

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

---

**Typographical conventions**

EMC uses the following type style conventions in this document:

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

---

**Where to get help**

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

**EMC product information**

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

**Technical support**

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

**Your comments**

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





# CHAPTER 1

## Introducing Data Domain Boost

This chapter contains the following topics:

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

## Revision History

The following table presents the revision history of this document.

**Table 1** Revision History of DD Boost Release 2.6

Revision	Date	Description
01	January 2015	Initial publication.

## Overview of DD Boost

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

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

---

### Note

A Data Domain system can be a single Data Domain system, a gateway, or a DD Extended Retention system.


---

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

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

The figure shows an example configuration of Data Domain Boost.

Figure 1 DD Boost — Configuration



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

## Supported Configurations

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

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

## Upgrade Compatibility

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

- All maintenance and patch versions within a *family* are backward compatible. A family is identified by the first two digits of the release number, such as 5.2. For example, 5.2.0.0, 5.2.0.2, 5.2.1.0, and 5.2.2.0 are all backward compatible.

- Replication is backward compatible across two consecutive release families, such as 5.5 and 5.4, although only the current release within each family is fully tested.
- Replication requires two systems: the destination system (the target) and the source system. The destination must be running the same version as, or one version newer than, the source.
- Both source and destination Data Domain systems must be licensed for replication.

# CHAPTER 2

## DD Boost Features

New and enhanced capabilities are available for Single Node and DD Extended Retention.

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

- [Overview of DD Boost Features](#)..... 14
- [Distributed Segment Processing](#) ..... 14
- [Managed File Replication](#)..... 15
- [IFGROUP: DD Boost IP Load Balancing and Failover](#)..... 16
- [DD Boost-over-Fibre Channel Transport](#)..... 19
- [DD Boost-over-Fibre Channel Path Management](#)..... 21
- [Virtual Synthetic Backups](#)..... 24
- [Backup Server Access Validation](#)..... 24
- [Storage Unit Capacity Quotas](#)..... 24

## Overview of DD Boost Features

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

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

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

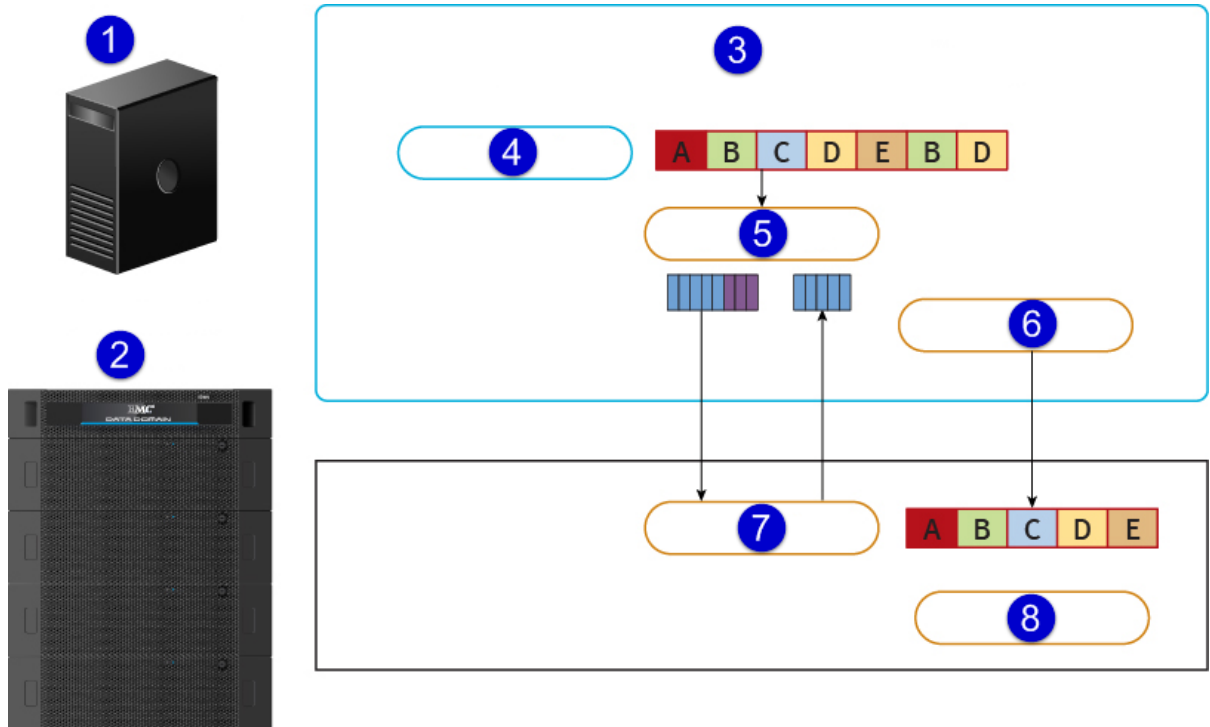
## Distributed Segment Processing

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 does not use extra memory, but some extra processing power on the application host is required to determine if the data is present on the Data Domain system.

Figure 2 Distributed Segment Processing Enabled



1. Database Server

2. Data Domain System
3. DD Boost Libraries
4. Segment
5. Fingerprint
6. Compress
7. Filter
8. Write

---

#### Note

When using the Solaris 11/11.1 bundled OpenSSL 1.0.0.j and running on either Solaris 11 (with SRU2 or higher) or Solaris 11.1 or higher, the DD Boost libraries offer improved distributed segment processing (DSP). DSP is enabled by default for Solaris environments with a SPARC T4 processor and running Solaris 11 (with SRU2 or higher) or Solaris 11.1 or higher.

---

## Managed File Replication

If supported by the backup application, 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.

## Low-Bandwidth Optimization

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

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

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

## Encrypted Managed File Replication

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

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

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

Encrypted managed file replication can be used with the encryption of data-at-rest feature available on the DD OS with the optional Encryption license. When encrypted

managed file replication is used with the encryption of data-at-rest feature, the encrypted backup image data is encrypted again using SSL for sending over WAN.

---

**Note**

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

## IFGROUP: DD Boost IP Load Balancing and Failover

---

**Note**

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

---

The IFGROUP 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.

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 IFGROUP 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.

---


**Note**

The managed file replication connection between the Data Domain systems is not part of IFGROUP. A single IP address is used for the destination Data Domain system. EMC recommends excluding one interface from the IFGROUP and reserving it for the managed file replication path between the source and destination Data Domain systems.

---



**Figure 3** IFGROUP: DD Boost IP Load Balancing and Failover



1. Backup Servers
2. DD Boost Libraries
3. Network Layer Aggregation
4. Load Balanced Backup Servers Group
5. Interface Group: Four Ports in the Group
6. NIC
7. DD Boost

IFGROUP can be used with other network layer aggregation and failover technologies. The sample configuration in the figure shows one possibility. The links connecting the backup servers and the switch connecting the Data Domain system are configured for an aggregated failover mode. This configuration provides network failover functionality from end-to-end. Any of the available aggregation technologies can be used between the backup server and the switch.

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

---

#### Note

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

---

IFGROUP 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 to 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.
- IFGROUP must be configured on Data Domain systems.
- Works with 1 GbE interfaces and 10 GbE interfaces in the same IFGROUP. Combining interfaces of different speeds in a single IFGROUP is allowed and supported.
- An administrator can define multiple IFGROUPs where load balancing and failover apply within an IFGROUP *<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 IFGROUP *<group-name>* is composed of a list of interfaces and clients that belong to the IFGROUP. Within an IFGROUP *<group-name>*, all interfaces are reachable by all the clients for *<group-name>*.
  - The IFGROUP *<group-name>* selection process checks the IFGROUP list in the following order, and selects the first match it finds.
    1. Client host-name (abc-11.d1.com)
    2. Client domain (\*.d1.com)
    3. All clients (\*)
    4. Client target IP address is an IFGROUP interface with a "dummy-name" configured to allow the IFGROUP to be enabled. Targeted IP selection allows clients to bypass the public IP address and connect directly to an IFGROUP interface.

## Interfaces

An ifgroup interface is a member of a single ifgroup *<group-name>* and may consist of the IPv4 associated with:

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

## Clients

An ifgroup client is a member of a single ifgroup *<group-name>* and may consist of:

- A fully qualified domain name (FQDN) such as `ddboost.datadomain.com`
- Wild cards such as `*.datadomain.com` or `**`
- A short name for the client, such as `ddboost`.

No client configuration is needed when the client connects directly to one of the interfaces of the ifgroup. If the client does not match the client criteria of an ifgroup, the last check is against the IP address being used as the administered interface for the Data

Domain system. This allows a "dummy" client name to be used to enable the ifgroup group.

---

**Note**

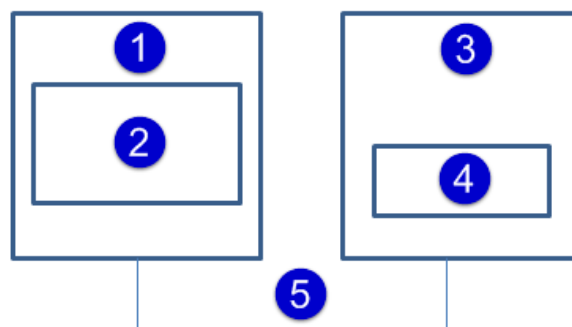
For this auto-detect client group to work, none of the other groups can use "\*" (all clients).

---

## DD Boost-over-Fibre Channel Transport

Prior to DD OS version 5.4, all communication between the DD Boost Library and any Data Domain system was performed using IP networking. The application specifies the Data Domain system using its hostname or IP address. See [Figure 4 on page 19](#).

**Figure 4** DD Boost-over-IP Transport



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

DD OS release 5.4 introduces an alternative transport mechanism for communication between the DD Boost Library and the Data Domain system — Fibre Channel.

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 5** SCSI Commands between Backup Server and Data Domain system.



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


Setting up the DD Boost-over-FC service on the Data Domain system requires additional configuration steps. See [Configuring DD Boost-over-FC Service on page 35](#) for details.

For the DD Boost-over-FC transport, load balancing and link-level high availability is achieved through a different means, not through IFGROUP. See the section [DD Boost-over-Fibre Channel Path Management on page 21](#) for a description.

**Note**

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

**Figure 6** Fibre Channel Communication Path



1. Backup Server
2. Application, DD Boost Library
3. IP or FC

4. IP or FC (Control)
5. Data Domain System, Replication Source
6. IP ONLY (Data)
7. Data Domain System, Replication Destination

## DD Boost-over-Fibre Channel Path Management


The IFGROUP-based mechanism described in [IFGROUP: DD Boost IP Load Balancing and Failover on page 16](#) is based on Ethernet interfaces and is not applicable to the Fibre Channel transport. Instead, a different path mechanism is provided for the DD Boost-over-FC solution.

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

For example, consider the case where:

- Backup server has 2 initiator HBA ports (A and B)
- Data Domain System has 2 FC target endpoints (C and D)
- Fibre Channel Fabric zoning is configured such that both initiator HBA ports can access both FC target endpoints
- Data Domain system is configured with a SCSI target access group containing:
  - Both FC target endpoints on the Data Domain System
  - Both initiator HBA ports
  - 4 devices (0, 1, 2, and 3)

**Figure 7** DD Boost-over-FC Path Management Scenario



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

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

- /dev/sg11: (A, C, 0)

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

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

## Initial Path Selection

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

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

## Dynamic Re-Balancing

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

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

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

## Client Path Failover

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

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

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

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

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

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

## Queue-Depth Constraints

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

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

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

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

Refer to [Sizing DD Boost-over FC Device-Set on page 36](#) for guidance regarding how many devices to configure based on the workload, type of Data Domain system, and whether or not the client system is queue-depth constrained.

## Virtual Synthetic Backups

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

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

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

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

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

## Backup Server Access Validation

Configuring backup server access validation for DD Boost limits access to the Data Domain system for DD Boost clients and removes dependency on the DNS. This feature enables a one time DD Boost authentication per connection. Connection authentication against the hostname is needed only until credentials are available. The list of clients can be updated at anytime without a restart requirement, thus eliminating access validation impact on jobs in progress.

## 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.



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

---

**Note**

Be careful with this feature when you are using backup applications (such as Symantec 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. So, logical quotas may need to be adjusted when the deduplication ratio changes.

---



# CHAPTER 3

## Preparing the Data Domain System for DD Boost

---

**Note**

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

---

This chapter covers the following topics:

- [Enabling DD Boost on a Data Domain System](#) ..... 28
- [Creating Storage Units](#) ..... 28
- [Configuring Quotas for Storage Units \(Optional\)](#) ..... 28
- [Configuring Distributed Segment Processing](#) ..... 29
- [Configuring IFGROUP](#) ..... 29
- [Configuring Managed File Replication \(MFR\)](#) ..... 33
- [Configuring Media Server Access Validation](#) ..... 34
- [Configuring DD Boost-over-FC Service](#) ..... 35
- [Sizing DD Boost-over FC Device-Set](#) ..... 36

## Enabling DD Boost on a Data Domain System

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

### Procedure

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

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

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

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

4. Set the DD Boost user by entering:

```
# ddbboost set user-name ddbboost-user
```

5. Enable DD Boost by entering:

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

### Note

- Only one user can be configured for DD Boost access on a Data Domain system at a time. The username, password, and role must have already been set up on the Data Domain system using the DD OS CLI 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
```

- The user must be configured in the backup application to connect to the Data Domain system.
- In DD OS 5.2 the user role of backup-operator is introduced. See the *EMC Data Domain Operating System 5.2 Administration Guide* for details.

## 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.

## Configuring 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 ddbboost in the *EMC Data Domain Operating System Command Reference Guide*.

### Procedure

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

```
# quota enable
```

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

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

3. To configure quota limits after a storage-unit has been created:

```
# quota set storage-units storage-unit-list
{soft-limit n {MiB|GiB|TiB|PiB}} {hard-limit n {MiB|GiB|TiB|PiB}}
```

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

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

## Configuring Distributed Segment Processing

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

The option can be configured using the following command:

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

---

### Note

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

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

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

## Configuring IFGROUP

---

### Note

This feature applies only to DD Boost over IP. For an overview of the IFGROUP feature, see [IFGROUP: DD Boost IP Load Balancing and Failover on page 16](#).

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

To configure IFGROUP, 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:

```
# ddboost ifgroup create group_name
```

**Examples:**

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

**Note**

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

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

```
# ddboost ifgroup add [group_name] {interface ipaddr | client host-name}
```

**Examples:**

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

**Note**

If no *group\_name* is specified, the default group is used.

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 ifgroup to register with the backup application. An interface that is not part of the ifgroup can also be used to register with the backup application.

EMC recommends that the interface be registered with a resolvable name using DNS or any other name resolution mechanism.

**Note**

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

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

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

**Note**

If no *group\_name* is specified, the default group is used.

**Note**

When clients connect directly to an ifgroup interface, client entry is not needed, but ifgroup requires one client to enable the group. Therefore, add a "dummy" client:

```
ddboost ifgroup add client group-name.noclient.
```

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

```
# ddbboost 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	ddbboost-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, `ddbboost-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.

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

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

---

**Note**

Removing the interface registered with the backup application makes the Data Domain system inaccessible to the backup servers. The configuration of the ifgroup 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 modify the interfaces, such as enabling an interface that is configured for ifgroup, execute the `ddboost show connections` command to update the load balancing view. Updating the load balancing view allows the ifgroup 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:

```
# ddboost ifgroup show connections
```

2. Ensure there are no pending jobs from backup servers connected to the Data Domain system.

3. Disable the *group-name* or default group on the system:

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

4. Reset the interface group:

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

**Results**

All the interfaces are released from the group. However, backup servers can still access the DD Boost storage server on the Data Domain system on the interface registered with the backup application. In the example above, the Data Domain system is still registered with the backup application using 192.168.1.1.

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

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

**Example:**

```
# ddboost 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 IFGROUP from the Enterprise Manager Data Management DD Boost view. (See the *EMC Data Domain Operating System Administration Guide*).

---

## Configuring Managed File Replication (MFR)

### Throttling MFR

If managed file replication is being used, replication throttling should be disabled. If throttling must be used, the workaround is to set the throttle manually with `/ddr/bin/repl_throttle` destination port rate command. Contact engineering if this is required.

### Enabling Low-Bandwidth Optimization

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

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

---

**Note**

Enabling or disabling the low-bandwidth optimization option does not require a restart of the Data Domain file system. However, after enabling low-bandwidth optimization, you need to run a full cleaning cycle on the Data Domain system for it to be effective.

---

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

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

---

**Note**

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

### Enabling Encryption

To enable the encrypted managed file replication option, enter:

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

The output indicates that the encryption you requested was enabled.

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

## Changing the MFR TCP Port

---

### Note

Changing the managed file replication TCP port requires a restart of the Data Domain file system. Therefore it should be a planned event.

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

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

---

### Note

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

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

## Configuring Media Server 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 add individual clients.

The backup server client list may contain both fully qualified domain names or short names. The client name must match the “hostname” on the backup host and is case sensitive.

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

```
# ddbboost access del clients *
```

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

```
# ddbboost access reset
```

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

```
# ddbboost access add clients <client-list>
```

Example:

```
# ddbboost access add clients ddbboost-dl.emc.com ddbboost-dlddbboost-dl.emc.com : Addedddbboost-dl : Added
```

To view the DD Boost client access list, enter:

```
# ddbboost access show
```

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

- Wild card \* followed by partial, for example, \*.emc.com followed by \*.com
- Perfect match of sent client name, for example, ddbboost-d1.emc.com

If the search is not found, the client will be denied access.

## Configuring DD Boost-over-FC Service

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

The following additional configuration steps are also required:

### 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>
```

Or accept the default, which is the base hostname of the Data Domain system. 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` endpoint, 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. Configure the device set of the SCSI target access group for one of the endpoint:

```
# ddboost fc group modify <group-name> device-set
count count endpoint endpoint-list
```

Example:

```
# ddboost fc group modify lab_group device-set count 8 endpoint 6a
```

6. Add initiators to the SCSI target access group:

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

Example:

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

## Sizing DD Boost-over FC Device-Set

As described in [DD Boost-over-Fibre Channel Path Management on page 21](#), 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 an initiator, by adding the initiator to a ddboost-type scsitarget access group:

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

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

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

The maximum number of supported DFC devices per Data Domain system is 64.

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

1. Is the backup server queue-depth constrained?  
As described in [DD Boost-over-Fibre Channel Path Management on page 21](#), Windows platforms are considered “queue-depth constrained,” because the Windows SCSI Pass-Through Interface mechanism will only conduct 1 SCSI request at a time through each of its generic SCSI devices. This impacts the performance of the DD Boost-over FC solution, if multiple connections (e.g. backup jobs) are trying to use the same generic SCSI device. So, for Windows platforms running more than one job, it is useful to advertise multiple DFC devices.

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

2. Number of physical paths between backup server and Data Domain system  
For each advertised DFC device, the backup server operating system will create  $n$  generic SCSI devices, one for each physical path through which the backup server OS can access the device.

For example, if:

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

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

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

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

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

## Sizing Calculation

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

### On the Data Domain System

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

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

Let C be the maximum number of connections per job:

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

Calculate:

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

### Example:

Assume:

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

Assume:

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

### Linux Backup Servers

The number of DFC devices advertised on the Data Domain system using the calculations listed under the [On the Data Domain System on page 37](#) section is sufficient for Linux backup servers. No additional configuration is required. Linux backup servers are not queue-depth constrained, so many connections can share the same DFC generic SCSI device with no performance impact.

### Windows Backup Servers

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

Let  $X$  = the number of DFC devices configured on the Data Domain system (from the [On the Data Domain System on page 37](#) section above). Let  $P$  = number of physical paths between backup server and Data Domain system. Let  $J$  = maximum number of simultaneous jobs, and let  $C$  = maximum number of connections per job:

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

Calculate:

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

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

Examples:

Assume:

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

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

Assume:

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

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

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

# CHAPTER 4

## Administration

---

**Note**

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

---

This chapter covers the following major topics:

- [Configuring a Backup Server](#) .....40
- [Backup Administration](#) .....40

## Configuring a Backup Server

Configure backup servers as specified by the backup application.

## Backup Administration

### Network Time-Outs

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

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

---

#### Note

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

---



# CHAPTER 5

## Basic Troubleshooting

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

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

This chapter covers the following topics:

- [General Troubleshooting](#)..... 42
- [Resolve time-out error](#)..... 42
- [Managed File Replication Job Fails](#).....42
- [Virtual Synthetic Backup](#)..... 43

## General Troubleshooting

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

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

---

### Note

A supported configuration can become unsupported if any component changes.

---

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

## 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
```

## Managed File Replication Job Fails

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

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

## Add license for Replication

### Procedure

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

```
# license add license code
```

## Verify Encrypted Managed File Replication Configuration

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

[Enabling Encryption on page 33](#) describes how to enable encrypted managed file replication.

## Virtual Synthetic Backup

- Verify that normal backups are OK.
- Verify that DDP\_SynWR RPCs are being sent.  

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

```
# ddbost option show
```

