



EMC[®] Data Domain[®] Boost for OpenStorage

Version 3.0

Administration Guide

302-001-271

REV. 02

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

Published April, 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.

EMC², 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

CONTENTS

	Preface	7
Chapter 1	Introducing Data Domain Boost for OpenStorage	11
	Revision History.....	12
	Overview of DD Boost for OpenStorage.....	12
	Supported Configurations	13
	Upgrade Compatibility.....	13
Chapter 2	DD Boost Features	15
	Overview of DD Boost Features.....	16
	Distributed Segment Processing	16
	Managed File Replication.....	17
	Low-Bandwidth Optimization.....	17
	Encrypted Managed File Replication.....	17
	Auto Image Replication (AIR).....	18
	Limitations When Using AIR on Data Domain Systems	19
	MTree Replication.....	19
	IPv6 Support.....	20
	IFGROUP: DD Boost IP Data Path Management.....	20
	Interfaces.....	21
	Clients.....	22
	IP Failover Hostname.....	23
	Interface Enforcement.....	24
	DD Boost-over-Fibre Channel Transport.....	25
	DD Boost-over-Fibre Channel Path Management.....	27
	Initial Path Selection.....	29
	Dynamic Re-Balancing.....	29
	Client Path Failover.....	29
	Queue-Depth Constraints.....	30
	Virtual Synthetic Backups.....	30
	Client Access Validation.....	31
	DD Boost Multiuser Data Path.....	31
	Storage Unit Management.....	31
	Multiuser Storage Units Access Control.....	31
	Storage Unit Capacity Quotas.....	32
	Storage Units Stream Count Management.....	32
Chapter 3	Preparing the Data Domain System for DD Boost	35
	Enabling DD Boost on a Data Domain System.....	36
	Assigning Multiple Users to DD Boost.....	36
	Creating Storage Units	37
	Configuring Logical Quotas for Storage Units (Optional)	38
	Configuring Storage Units with Stream Limits (Optional).....	38
	Configuring Distributed Segment Processing.....	39
	Configuring IFGROUP	40
	Modifying an Interface Group.....	42
	Removing an Interface Group.....	43

	Configuring Managed File Replication (MFR)	43
	Throttling MFR	44
	Enabling Low-Bandwidth Optimization	44
	Enabling Encryption	44
	Enabling IPv6 Support	44
	Changing the MFR TCP Port	45
	Configuring Client Access Validation	45
	DD Boost Integration with Secure Multi-Tenancy	46
	Configuring DD Boost-over-FC Service	47
	Sizing DD Boost-over FC Device-Set	48
	Sizing Calculation	49
Chapter 4	Installing DD Boost for OpenStorage	53
	Installation Overview	54
	OST Plug-In and DD OS Upgrades	54
	Firewalls and Ports	54
	Installing OST Plug-In for NetBackup	55
	Installing the OST Plug-In on Media Servers	55
	Installing the UNIX Plug-In	55
	Installing the Windows Plug-In	56
	NetBackup Services	57
	Installing OST Plug-In for Backup Exec	58
	Installing the Plug-In on Media Servers	58
	Install the Windows Plug-In	58
	Backup Exec Services	59
	Tuning Windows Media Servers for Performance	59
	Uninstalling the Windows Plug-in	59
Chapter 5	Backup Application Administration	61
	Configuring a Media Server	62
	NetBackup Configuration	62
	Backup Exec Configuration	72
	NetBackup Administration	76
	Find your OST Plug-in Version	76
	Find your NetBackup version	76
	Network Time-Outs	76
	Grouping Storage Units to Provide Failover	77
	Backup Exec Administration	78
	Find your OST plug-in version	78
	Find your Backup Exec version	78
	Delete Storage Units on Data Domain Systems	78
Chapter 6	Basic Troubleshooting	79
	General Troubleshooting	80
	Data Domain System Settings for File Replication	80
	NetBackup Troubleshooting	80
	Unable to Delete the Data Domain System	80
	Check the Installation	81
	Check Credentials	82
	Resolve License Errors	82
	Error Logging on the Media Servers	82
	Resolving Failed Backups on Media Servers	82
	Resolve Plug-In Log Messages	84

Resolve “Cannot connect on socket” Error.....	84
NetBackup Backup Jobs Fail on Solaris Media Servers.....	84
Optimized Duplication Job Fails.....	85
Virtual Synthetic Backup.....	85
Monitoring Auto Image Replication.....	86
Backup Exec Troubleshooting.....	91
Basic Troubleshooting.....	91
Check the installation.....	91
Check Credentials for a Data Domain System.....	91
Resolve License Errors	91
Set Up Active Debugging.....	91

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 for OpenStorage 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 Symantec 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 for OpenStorage Backup Application Documentation

Documentation for backup applications is available through the Symantec web site.

Symantec NetBackup Documentation

From the general Symantec support page, navigate to the NetBackup Server product page and search the knowledge base for Documentation.

Note

To locate a document, enter its title as a search criterion in your favorite search engine.

- *NetBackup Shared Storage Guide*
- *NetBackup Troubleshooting Guide*

- *NetBackup Commands*

See these NetBackup documents for more information:

- *NetBackup Backup, Archive, and Restore Getting Started Guide*
- *NetBackup Administrator's Guide for UNIX and Linux* (two volumes)
- *NetBackup Administrator's Guide for Windows* (two volumes)
- *Best Practices for using Storage Lifecycle Policies in NetBackup.*
- *NetBackup 7.x Hardware Compatibility List* that includes information for supported OpenStorage servers.

Symantec Backup Exec Documentation

This document is installed with the application:

- *Symantec Backup Exec 2012 Administrator's Guide*
- *Backup Exec 2012 Hardware Compatibility List* that includes information for supported OpenStorage servers.

Special notice conventions used in this document

EMC uses the following conventions for special notices:

NOTICE

Identifies content that warns of potential business or data loss.

Note

Contains information that is incidental, but not essential, to the topic.

Typographical conventions

EMC uses the following type style conventions in this document:

Table 1 Typography

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

Where to get help

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.

CHAPTER 1

Introducing Data Domain Boost for OpenStorage

This chapter contains the following topics:

- [Revision History](#)..... 12
- [Overview of DD Boost for OpenStorage](#)..... 12
- [Supported Configurations](#)13
- [Upgrade Compatibility](#).....13

Revision History

The following table presents the revision history of this document.

Table 2 Revision History of DD Boost for OpenStorage Administration Guide, Release 3.0.x

Revision	Date	Description
02 (3.0.2)	April 20, 2015	<ul style="list-style-type: none"> Documented several IFGROUP enhancements: <ul style="list-style-type: none"> Full support for static IPv6 addresses, providing the same capabilities for IPv6 as for IPv4. A client selection option that uses a range of IP addresses to route inbound clients from public to private networks. Ability to enforce private network connectivity, ensuring that a failed job does not reconnect on the public network. Added HP-UX as a supported client operating system for the DD Boost-over-Fibre Channel feature.
01 (3.0.1)	October 14, 2014	<ul style="list-style-type: none"> Reorganized the content to better reflect the sequence of tasks involved in working with DD Boost for OST. Added information about the new Failover Hostname feature. Removed TCP 3008 (RSS) from the list of required ports. Mentioned that DSP is enabled by default on Solaris plug-ins running on a SPARC T4 processor.

Overview of DD Boost for OpenStorage

In the context of Symantec 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.


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.

Figure 1 DD Boost for OpenStorage — NetBackup Configuration



1. Clients
2. Server
3. Primary Storage
4. Media Server
5. OST Plug-in
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 OST plug-in version 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](#)..... 16
- [Distributed Segment Processing](#) 16
- [Managed File Replication](#)..... 17
- [Auto Image Replication \(AIR\)](#)..... 18
- [MTree Replication](#)..... 19
- [IPv6 Support](#)..... 20
- [IFGROUP: DD Boost IP Data Path Management](#).....20
- [DD Boost-over-Fibre Channel Transport](#)..... 25
- [DD Boost-over-Fibre Channel Path Management](#).....27
- [Virtual Synthetic Backups](#).....30
- [Client Access Validation](#)..... 31
- [DD Boost Multiuser Data Path](#)..... 31
- [Storage Unit Management](#)..... 31

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.

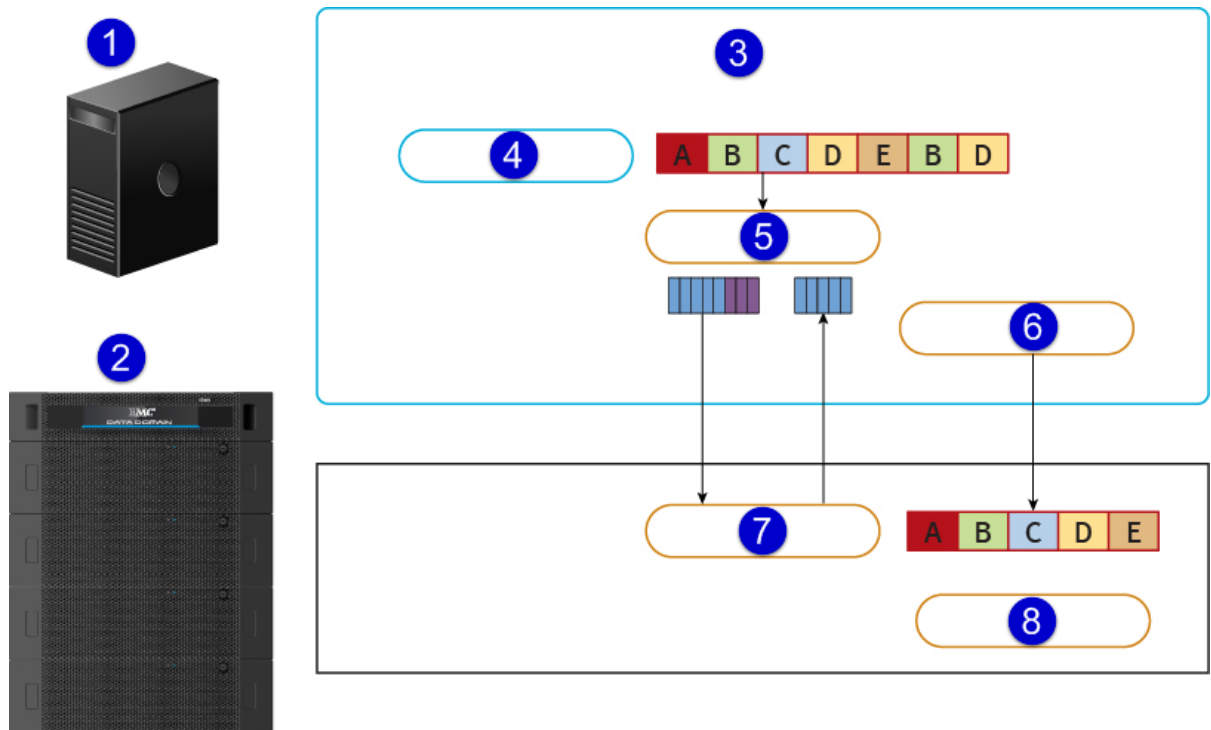
Distributed Segment Processing

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. 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 higher) or Solaris 11.1 or higher, the plug-in offers improved distributed segment processing (DSP). DSP is enabled by default for Solaris plug-ins running on a SPARC T4 processor and running Solaris 11 (with SRU2 or higher) or Solaris 11.1 or higher.

Managed File Replication

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.
-

Auto Image Replication (AIR)

Auto Image Replication (AIR) works by duplicating images to a remote master server domain. The AIR feature, introduced in NetBackup 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 duplication using AIR. Electronic off-siting in this manner allows the backup set to be duplicated 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 duplicate copy is available at the disaster recovery site as soon as the duplication has completed.

In order 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 figure below.

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 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 duplicate the backup to other storage locations in the target domain.

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.

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

MTree Replication

Beginning with DD OS Release 5.5, MTree replication for storage units is supported for different user names on source and destination Data Domain systems. To enable MTree replication for storage units, you must convert the target storage unit from MTree to storage unit by assigning a user to the MTree. To assign a user to the MTree, use the DD

OS `ddboost storage-unit modify` command. (See the *EMC Data Domain Operating System Command Reference Guide* for details.)

IPv6 Support

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

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

IFGROUP: DD Boost IP Data Path Management

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 (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 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.

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 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 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>* includes 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>*.
- 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 22.

Interfaces


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

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

IFGROUP provides full support for static IPv6 addresses, providing the same capabilities for IPv6 as for IPv4. Concurrent IPv4 and IPv6 client connections are allowed. A client connected with IPv6 sees IPv6 IFGROUP interfaces only. A client connected with IPv4 sees

IPv4 IFGROUP interfaces only. Individual IFGROUPs include all IPv4 addresses or all IPv6 addresses.

Figure 4 IFGROUP Support for IPv4 and IPv6 Addressing



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`
- Client public IP range, such as `128.5.20.0/24`

Prior to write or read processing, the client requests an IFGROUP IP address from the server. To select the client IFGROUP association, the client information is evaluated according to the following order of precedence (see [Figure 5 on page 23](#)):

1. Client Name: `abc-11.d1.com`
2. Client Domain Name: `*.d1.com`
3. All Clients: `*`
4. IP address of the connected Data Domain system. If there is already an active connection between the client and the Data Domain system, and the connection exists on the interface in the IFGROUP, then the IFGROUP interfaces are made available for the client.
5. Connected client IP range. An IP mask check is done against the client source IP; if the client's source IP address matches the mask in the IFGROUP clients list, then the IFGROUP interfaces are made available for the client.
 - For IPv4, `xx.xx.xx.0/24` (`128.5.20.0/24` in [Figure 5 on page 23](#)) provides a 24-bit mask against the connecting IP. The /24 represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.
 - For IPv6, `xxxx::0/112` provides a 112-bit mask against the connecting IP. The /112 represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.

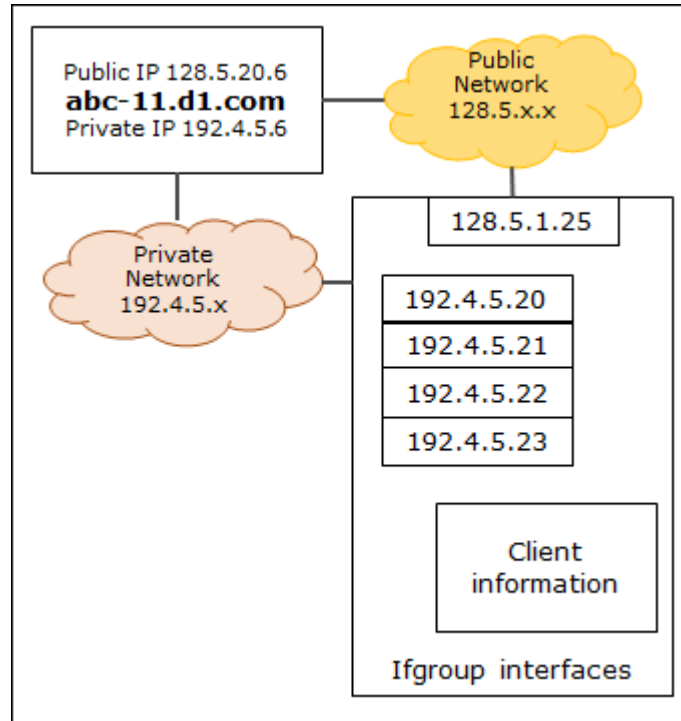
This host-range check is useful for separate VLANs with many clients where there isn't a unique partial hostname (domain).

Note

In a mixed network with IPv4 and IPv6 addressing, IFGROUP 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 (*.emc.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, IFGROUP interfaces are not used for this client.

Figure 5 IFGROUP Host Range for Client Selection



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 ifgroup so that you can connect directly to the ifgroup 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, ifgroup 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.

Interface Enforcement

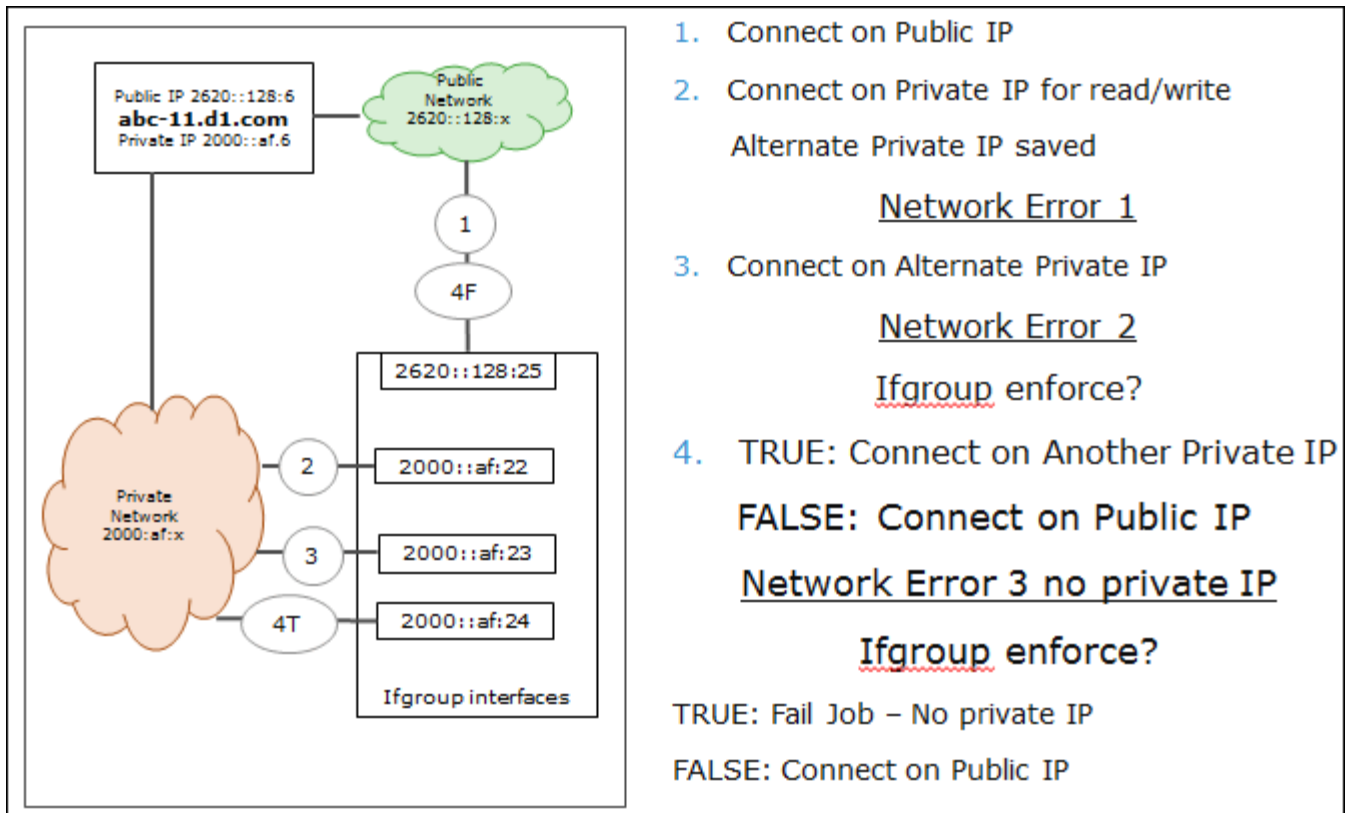
IFGROUP gives you the ability to enforce private network connectivity, ensuring that a failed job does not reconnect on the public network after network errors. When interface enforcement is enabled, a failed job can only retry on an alternative private network IP address. Interface enforcement is only available for clients that use IFGROUP interfaces.

Interface enforcement is off (FALSE) by default. To enable interface enforcement, you must add the following setting to the system registry:

`system.ENFORCE_IFGROUP_RW=TRUE`


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

[Figure 6 on page 25](#) shows the decision flow for IFGROUP 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 6 IFGROUP Connection Decision

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 7 on page 25](#).

Figure 7 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, and HP-UX 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 8 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 on page 47](#) 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 27](#) 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 9 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 IFGROUP-based mechanism described in [IFGROUP: DD Boost IP Load Balancing and Failover on page 20](#) 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 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 10 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 48](#) 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.

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 after each restart of DD Boost (Enable/Disable), file system restart, or Data Domain system reboot. 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.

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 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 system user for each backup application (NetBackup or Backup Exec) to be used for their storage units. The storage units are either created with a username, or can be modified to change the username for an upgrade. 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 Secure Multi-Tenancy feature. One storage unit must be configured for each tenant unit. Each tenant unit can be associated with multiple storage units. Tenant unit association and storage unit username ownership are independent from each other. The tenant unit is used for management path using the command-line-interface, but cannot be used for data path, for example, read and write. All commands for storage units support tenant units.

Note

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

Storage Unit Capacity Quotas

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

Optionally, you can configure the reported physical size. The physical size is the Disk Pool "raw size" in NetBackup. 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`.

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.

Storage Units Stream Count Management

This feature adds stream count management to storage units. The output of the `ddboost show connections` command provides information about the stream counts based on connections associated with the stream types on the Data Domain system. This feature adds per storage unit stream count information and the ability to configure warning limits per stream type to generate alerts.

The system administrator has the option of configuring stream warning limits for each storage unit for each of the four stream types: backup, restore, replication, and combined streams. Any of these stream warning limits can also be set to `none`. For each storage unit, four stream counters are maintained to monitor backup, restore, replication-in, replication-out data. Warning limits can be configured for the stream types using the four new keywords added to the `ddboost storage-unit create` command for the backup, restore, replication, and combined streams.

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 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 38](#).

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 *EMC Data Domain Operating System Command Reference Guide*.

This chapter covers the following topics:

• Enabling DD Boost on a Data Domain System	36
• Assigning Multiple Users to DD Boost	36
• Creating Storage Units	37
• Configuring Logical Quotas for Storage Units (Optional)	38
• Configuring Storage Units with Stream Limits (Optional)	38
• Configuring Distributed Segment Processing	39
• Configuring IFGROUP	40
• Configuring Managed File Replication (MFR)	43
• Configuring Client Access Validation	45
• DD Boost Integration with Secure Multi-Tenancy	46
• Configuring DD Boost-over-FC Service	47
• Sizing DD Boost-over FC Device-Set	48

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 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:

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

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

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

4. Enable DD Boost by entering:

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

Note

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

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

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

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

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

```
# ddbboost user assign jsmith
```

Assigning Multiple Users to DD Boost

As system administrator, you need to create a local Data Domain system user for each backup application to use with their storage units. The storage units are either created with a username, or can be modified to change the username for an upgrade. Storage units are accessible only to applications with the username that owns the storage unit. Each storage unit is owned by one username, and the same username can own multiple storage units. The application 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 `ddbboost user assign` command. If a storage

unit is created without specifying the owning username, the current DD Boost user name is assigned as owner.

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

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

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

```
# ddbboost user show
DD Boost user
-----
user1
user2
-----
```

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

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

Note

The `ddbboost file-replication show` commands have been updated to filter information for the storage units.

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:

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

Note

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

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

Note

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

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

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

Note

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

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

```
# ddbboost storage-unit show
```

Name	Pre-Comp (GiB)	Status	User
backup	19517.4	RW	user1
tset	0.0	RW	user2
SYNTH_REPL	2221.0	RW	user3

```

-----
Q      : Quota Defined
RO     : Read Only
RW     : Read Write

```

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 `ddbboost` in the *EMC 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 85210), 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 four limits: backup, restore, replication and combined streams. When any

stream count exceeds the warning limit quota, an alert is generated. The alert automatically clears once the stream limit returns below the quota for over 10 minutes.

Note

DD Boost users are expected to reduce the workload to remain below the stream warning quotas or the system administrator can change the warning limit configured to avoid exceeding the limit.

To create a storage unit with stream limits, enter:

```
# ddbost storage-unit create NEW_STU0 user user2 write-stream-soft-limit 5
read-stream-soft-limit 1 repl-stream-soft-limit 2
Created storage-unit "NEW_STU0" for "user2".
Set stream warning limits for storage-unit "NEW_STU0".
```

To modify the stream limits for storage units, enter:

```
# ddbost storage-unit modify NEW_STU1 write-stream-soft-limit 3
read-stream-soft-limit 2 repl-stream-soft-limit 1
NEW_STU1: Stream soft limits: write=3, read=2, repl=1, combined=none
```

To display the DD Boost stream limits for all the active storage units, enter:

```
# ddbost streams show active
```

Name	Read Streams	Write Streams	Repl-out Streams	Repl-in Streams	Read Limit	Write Limit	Repl Limit	Combined Limit
NEW_STU1	1	0	0	0	2	3	1	
NEW_STU0	0	2	1	0	1	5	2	

```
DD System Stream Limits: read=50 write=180 repl-in=180 repl-out=90 combined=180
```

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:

```
# ddbost 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 YYYY/MM/DD	Time HH:MM	read streams	write streams	repl-out streams	repl-in 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.

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 20](#).

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:

```
# ddbboost ifgroup create group_name
```

Examples:

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

Note

The *group_name* “default” can be used without being created first. In all the remaining `ddbboost 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.

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

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

Examples:

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



```
# ddboost ifgroup add interface 10.6.109.141 client *

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

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

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

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. 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 ifgroup:

```
# ddboost 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:

```
# ddboost 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:

```
# ddboost 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.

```
# ddboost 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 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 media 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, media 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 System 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 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 `ddbboost 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:

```
# ddbboost 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:

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

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

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

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 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 add individual clients.

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

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 {medium | high}]
```

Example:

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

To view the DD Boost clients list, enter:

```
# ddbboost clients show config
```

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

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

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

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

DD Boost Integration with Secure Multi-Tenancy

The tenant-unit keyword is introduced to the `ddbboost storage-unit` command for the DD boost integration with Secure Multi-Tenancy (SMT). One tenant unit must be configured for each storage unit and 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 command, but cannot be used for data path for example, read and write. All commands related for storage units support tenant units.

To create a SMT tenant unit to integrate with the DD Boost user, enter:

```
# smt tenant-unit create tu1Tenant-unit
"tu1" created.
```

To assign a DD Boost user to a default tenant unit, enter:

```
# ddbboost user option set user1 default-tenant-unit tu1
Default-tenant-unit is set to "tu1" for user "user1".
```

Note

Assign more DD Boost users to their default tenant units.

To delete and unassign a DD Boost user from the default tenant unit, enter:

```
# ddbboost user option reset user3
Default-tenant-unit is reset for user "user3".
```

To display the DD Boost user associated with a default tenant unit, enter:

```
# ddbboost user show
DD Boost user   Default tenant-unit
-----
user1           tu1
user2           tu2
```

Configuring DD Boost-over-FC Service

Before you begin

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

Note

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

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:

```
# ddboost option set fc enabled
```

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

```
# ddboost 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:

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

Example:

```
# ddboost 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:

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

Example:

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

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

```
# scsitarget initiator show list
```

7. Add initiators to the SCSI target access group:

```
# 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 27](#), 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 media server? The answer depends upon several factors:

1. Is the media server queue-depth constrained?
As described in [DD Boost-over-Fibre Channel Path Management on page 27](#), 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 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. EMC recommends that the same number of DFC devices be advertised 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 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 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 on page 49](#) 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 on page 49](#)). 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 on page 49](#) 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`.

CHAPTER 4

Installing DD Boost for OpenStorage

Note

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

This chapter covers the following topics:

- [Installation Overview](#) 54
- [Installing OST Plug-In for NetBackup](#) 55
- [Installing OST Plug-In for Backup Exec](#) 58
- [Tuning Windows Media Servers for Performance](#) 59
- [Uninstalling the Windows Plug-in](#) 59

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. You can purchase a DD Boost license key directly from EMC.
 - 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 user-name per storage-unit or one user-name for multiple storage-units.
3. Install the OST plug-in software on each media server.
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 35](#).

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*. The OST plug-in and the DD OS compatibility is defined in the *Data Domain Boost Compatibility Guide*.

To take advantage of new features in a DD OS release, 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 on page 55](#).

Note

This document illustrates the DD Boost configuration on Data Domain using commands in DD OS 5.5. 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 EMC 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 -d directorypath
```

The directory path is optional. If you do not specify a directory path, the script uses `/usr/opensv/lib/ost-plugins`, which is where the backup application normally looks for packages. The shared library files that the script installs are `libstspiDataDomain.so` and `libstspiDataDomainMT.so`.

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
```

Correcting a Failure in UNIX Plug-In Installation or Update (AIX Media Servers)

Procedure

1. Stop the NetBackup Remote Manager and Monitor Service (`nbrmms`).
2. Enter:

```
# ./install.sh
a plugin already exists, do you want to proceed with installation?
(y or n) y
Installing the Data Domain OpenStorage Client Libraries ...
cp libstspiDataDomain.so /usr/opensv/lib/ost-plugins/
libstspiDataDomain.socp:
/usr/opensv/lib/ost-plugins/libstspiDataDomain.so:
Cannot open or remove a file containing a running program.
ERROR in copying libstspiDataDomain.so to
/usr/opensv/lib/ost-plugins/libstspiDataDomain.so, error = 1
```

The install script might fail and display an error message that indicates that the plug-in already exists. This failure occurs if a plug-in is already installed and is being replaced by another instance.

3. If the install script fails:

- a. Enter:

```
# /usr/sbin/slibclean
```

- b. Enter:

```
# ./install.sh
```

Because the modules have now been unloaded from memory, the `install.sh` script should now run correctly.

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 the EMC 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 on page 57](#) 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 57](#).
6. Tune the Windows media server for performance. See [Tuning Windows Media Servers for Performance on page 59](#).

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 the EMC Online Support.
 - b. Stop any Backup Exec services. Follow the instructions given in [Backup Exec Services on page 59](#) 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 on page 59](#).
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 85209, which is available on the EMC Support portal site <https://support.emc.com>.

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 Window's 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 *EMC Data Domain Operating System Command Reference Guide*.

This chapter covers the following major topics:

- [Configuring a Media Server](#).....62
- [NetBackup Administration](#)..... 76
- [Backup Exec Administration](#)..... 78

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.x
 - 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 on page 57](#).
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 7 allows 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:

```
# bpstuaadd -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.

EMC recommends 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. EMC recommends that you 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 71960, which is available on the EMC 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 Symantec 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.

```
# ddboost 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.

```
# 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

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.

```
# 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-
```

```

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 `ddboost show stats` and `ddboost show histograms` commands.

```
# ddboost show stats
07/06 07:13:38
```

DD Boost statistics:

```

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

```

	Count	Errors
DDP_SYNWR	18	0

-----	-----	-----
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.

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 *Symantec Backup Exec 2012 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 `ddboost 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: $48 \geq \# \text{ OpenStorage storage units} + \Sigma \text{ concurrent operations for each storage unit}$ 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 on page 64](#).
2. From the Backup Exec Home page, select **Configure Devices** from the **Tools** menu. Select **Add OpenStorage** from the menu.
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 *Symantec Backup Exec 2012 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.

EMC recommends that you add the destination Data Domain system's IP address to the source Data Domain system using the 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 Symantec 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 7.5
```

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.

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.

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:

```
# bpexptime -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 -deletedp -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 on page 80](#).

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:

```
# ddboost 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 which is available at <https://support.emc.com>.

This chapter covers the following topics:

- [General Troubleshooting](#)..... 80
- [Data Domain System Settings for File Replication](#)..... 80
- [NetBackup Troubleshooting](#)..... 80
- [Backup Exec Troubleshooting](#)..... 91

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. [Configuring a Media Server on page 62](#) 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 EMC 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 page 82](#).

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 on page 82](#).

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/opensv/lib/ost-plugins/
```

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

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

- 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`.

2. 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

1. 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 on page 63](#), 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

1. Open the file `bp.conf`.
2. Check if it contains an extra `CLIENT_NAME` entry.
3. 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

1. 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:

```
# filesys 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 on page 64](#).
 - **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:

```
# filesys 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 apodrrp01: 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> apodrrp01: /usr/opensv/lib/ost-plugins/
libstspiDataDomain.so:stspi_open_server STS_EPLUGIN Can't connect to
mountd on apodrrp01 (: 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
```

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:

```
# ddboost association show
```

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

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

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

```
# ddboost 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:

```
ddboost 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 -updatests -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 -updatests -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:sparcl2sol02 )
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: sol022sparcl1
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:sol022sparcl1 )
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, look for Replication Target and Source, enter:

```
nbdevquery -listdp -stype DataDomain -U
Disk Pool Name      : sol022sparcl-dd670c2
Disk Pool Id       : sol022sparcl-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-sparcl.datadomain.com_ddr1.domain1.com_1328637954_1.imgset
```

Where: `192:rtp-ost-sparcl.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:

```
bluemedi.datadomain.com_31234_6589_1.event.000000000000000006
```

Where: `bluemedi.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 000000000000000006` is the unique event identifier on this Data Domain system.

- The long term presence (more than 2 hours) of `.imgset` files in the target storage unit indicates that the target NetBackup domain is not querying for posted events.
- The long term presence (more than 2 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-dd670c2.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-sparcl.datadomain.com> media=<rtp-ost-
sparcl.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 `ddboost 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, 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, `bluemia_1329923106`.

```
07:05:29.624 [24145] <2> process_requests: fork cmd =
/usr/opensv/netbackup/bin/bpdm bpdm -restoretir -cmd -b
bluemia_1329923106 -c
bluemia -cn 1 -drn 0 -p @aaaal -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
bluemia_1329923106
```

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

```
7:05:30.232 [24150] <2> bpdm: INITIATING (VERBOSE =
5): -restoretir -cmd -b bluemia_1329923106 -c
bluemia -cn 1 -drn 0 -p @aaaal -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 ostga-
sparcl.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. In case, 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 Symantec Best Practices Guide, enter:

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

Note

See the Symantec NetBackup Auto Image Replication FAQ at <http://www.symantec.com/business/support/>.

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 74](#).
- 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\`.
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.

