



Data Domain Boost for OpenStorage Administration Guide

2.5

Backup Recovery Systems Division
Data Domain LLC
2421 Mission College Boulevard, Santa Clara, CA 95054
866-WE-DDUPE; 408-980-4800

May 4, 2012
759-0017-0002 Revision A

Copyright © 2012-2012 EMC Corporation. All Rights Reserved.

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC, Data Domain, and Global Compression are registered trademarks or trademarks of EMC Corporation in the United States and/or other countries.

All other trademarks used herein are the property of their respective owners.

Contents

- About This Guide 7
 - Audience 7
 - Related Documents 7
 - DD Boost for OpenStorage Backup Application Documentation 7
 - Access to Documents at Data Domain 9
 - Access Data Domain Documents 9
 - Conventions 10
 - Data Domain Support 10
- 1 Introducing Data Domain Boost for OpenStorage . . . 11
 - Supported Configurations 12
- 2 DD Boost Features 13
 - Distributed Segment Processing 14
 - Managed File Replication 16
 - Low-Bandwidth Optimization 17
 - Encrypted Managed File Replication 18
 - Advanced Load Balancing and Link Failover 19
 - Interfaces 22
 - Clients 23
 - Heterogeneous Application Host Support 23
 - Virtual Synthetic Backups 23

Storage Quotas	24
Media Server Access Validation	25
3 Installing DD Boost for OpenStorage	27
Installation Overview	27
OST Plug-In and DD OS Upgrades	28
Firewalls and Ports	28
Enabling DD Boost on a Data Domain System	29
NetBackup Installation	30
Installing the OST Plug-In on Media Servers	31
NetBackup Services	34
NetBackup UNIX Services	34
NetBackup Windows Services	34
Backup Exec Installation	35
Installing the Plug-In on Media Servers	35
Backup Exec Services	37
Tuning Windows Media Servers for Performance	38
Uninstalling the Windows Plug-in	38
4 DD Boost for OpenStorage Configuration	39
Configuring the Data Domain System	39
Creating Storage Units	40
Configuring Quotas for Storage Units (Optional)	40
Configuring Distributed Segment Processing	41
Configuring Advanced Load Balancing and Link Failover	42
Enabling Low-Bandwidth Optimization	46
Enabling Encrypted Optimized Duplication	47

Configuring Media Server Access Validation	48
Configuring a Media Server	49
NetBackup Configuration	49
Concurrent Connection Limit	50
Registering Each Data Domain System	50
Adding Credentials	51
Creating Disk Pools	51
Creating Storage Units	55
Creating a Backup Policy	58
Configuring Buffers	58
Configuring Optimized Duplication	60
Using Storage Lifecycle Policies to Automate Optimized Duplication	61
Configuring a Virtual Synthetic Backup	65
Sample Backup Operations	67
Backup Exec Configuration	71
Creating a Logon Account	72
Configuring Devices	72
Configuring Optimized Duplication	75
5 Administration	89
Data Domain System Administration	89
Restricting Export Permissions to DD Boost	89
Optimized Duplication Version Compatibility	91
Modifying an Interface Group	91
Removing Advanced Load Balancing and Link Failover Configuration.	93

NetBackup Administration	94
Network Time-Outs	94
Grouping Storage Units to Provide Failover	95
Backup Exec Administration	96
6 Basic Troubleshooting	99
General Troubleshooting	99
Data Domain System Settings for File-Replication	100
NetBackup Troubleshooting	100
Unable to Delete the Data Domain System	100
Error Logging on the Media Servers	103
Resolving Failed Backups on Media Servers	104
NetBackup Backup Jobs Fail on Solaris Media Servers . .	107
Optimized Duplication Job Fails	108
Virtual Synthetic Backup	108
Backup Exec Troubleshooting	109
Basic Troubleshooting	109

About This Guide

This guide explains how to install, configure, and use Data Domain Boost for OpenStorage 2.5 with Data Domain systems.

Audience

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

Related Documents

The following Data Domain system documents provide additional information:

- *DD OS Release Notes Version 5.2*
- *DD OS 5.2 Initial Configuration Guide*
- *DD OS 5.2 Administration Guide*
- *DD OS 5.2 Command Reference Guide*
- *Data Domain Expansion Shelf Hardware Guide*
- The Data Domain system installation and setup guides 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 6.5.3 and 6.5.4.*
- *NetBackup 6.5 Hardware Compatibility List* that includes information for supported OpenStorage servers.
- *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 2010 Administrator's Guide*
- *Backup Exec 2010 Hardware Compatibility List* that includes information for supported OpenStorage servers.

Access to Documents at Data Domain

The Documentation page at <https://my.datadomain.com/documentation> provides access to documents that are related to use of Data Domain products:

- User guides, under Product Documentation.
- Guides for how to integrate Data Domain systems with backup applications, under Integration Documentation.
- Matrices that show which components are compatible with each other, under Compatibility Matrices:
 - Data Domain hardware product numbers
 - Data Domain operating system (DD OS) versions
 - Backup software versions
 - Backup software server and client operating system versions

Access Data Domain Documents

1. Log into the support portal at:
<https://my.datadomain.com/documentation>.
2. To view end user documents, click Product Documentation and then perform the following steps:
 - a. Select your product (for example, DD Boost OpenStorage).
 - b. Use the pull-down menu to select your Platform (for example, Windows).
 - c. Click View then select the desired title.
3. To view integration-related documents, perform the following steps:
 - a. Click Integration Documentation.
 - b. Select a vendor from the Vendor menu.
 - c. Select the desired title from the list and click View.
4. To view compatibility matrices, perform the following steps.

- a. Click Compatibility Matrices.
- b. Select the desired title from product menu and click View.

Conventions

The following tables describe the typographical conventions used in this guide. The following table describes the typographical conventions used in this guide.

Typeface or Symbol	Usage	Examples
Monospace	Commands, command options, and parameters and computer output.	Use the <code>config</code> command to manage the Data Domain system configuration settings.
Monospace bold	Commands the user types at the command prompt (#).	Enter: # config setup
<i>Monospace italic bold</i>	Command variables the user types at the command prompt (#).	# log view <i>file_name</i>
<i>Italic</i>	Book titles, and variables.	Refer to the <i>DD OS Command Reference Guide</i> for complete descriptions of DD OS commands.
Pipe () and curly braces ({})	Choose (pipe) between a required argument (curly braces) in the CLI.	{ <i>arg1</i> <i>arg2</i> }
Brackets ([]) and ellipses (...)	One or more (list with commas and ellipses) optional (bracket) arguments in the CLI.	[<i>arg1</i> , <i>arg2</i> , ...]

Data Domain Support

To resolve issues with Data Domain products, contact your contracted support provider or visit us online at <https://my.datadomain.com>.

1 Introducing Data Domain Boost for OpenStorage

Data Domain Boost (DD Boost) for OpenStorage enables media servers to communicate with storage systems without the need for Data Domain storage systems to emulate tape. In the context of Symantec backup applications (NetBackup and Backup Exec), there are two components to the software:

- An OST plug-in that you install on each media server. This plug-in includes the DD Boost libraries for integrating with the DD Boost server running on the Data Domain system.
- The DD Boost server that runs on Data Domain systems. A Data Domain system can be a single Data Domain system, a gateway, or a DD Extended Retention system or a Global Deduplication Array (GDA).

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 application manages all files (collections of data) in the catalog, even those created by the Data Domain system.

The Data Domain system exposes pre-made disk volumes called storage units to a DD Boost-enabled 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 it is supported by Data Domain.

Figure 1-1 shows an example configuration of Data Domain Boost for Open Storage using NetBackup.

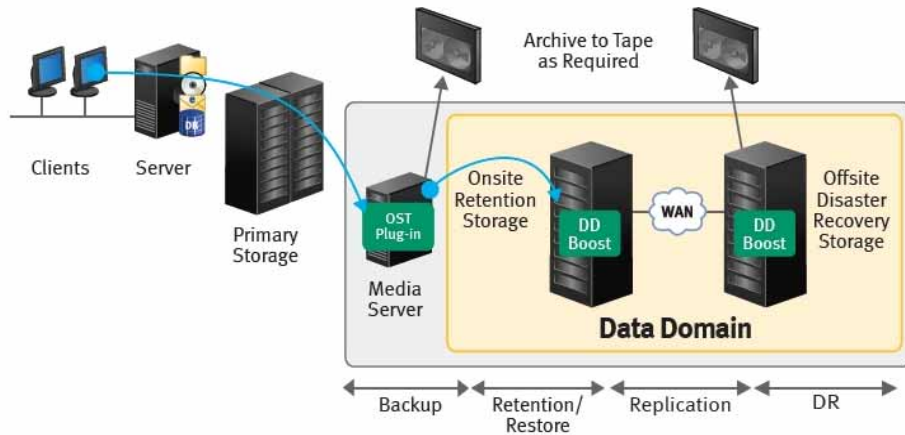


Figure 1-1: DD Boost for OpenStorage - NetBackup Configuration

Supported Configurations

Data Domain supports DD Boost on the following Data Domain systems:

- Appliance Series: DD990, DD890, DD880, DD860, DD690, DD670, DD660, DD640, DD630, DD620, DD610, DD580, DD565, DD530, DD510, DD160, DD140, DD120
- Gateway Series: DD880g, DD690g, DD580g
- Global Duplication Array Series: Based on DD 890 and DD880
- DDX Array Series: DD880, DD690, DD580, DD880g, DD690g, DD580g

The OST plug-in version must be compatible with your Data Domain system and backup application configurations. Data Domain does not support combinations other than those detailed in the *DD Boost for OpenStorage Compatibility List*, which is available from the Data Domain Support portal <https://my.datadomain.com>. Navigate to the Documentation > Compatibility Matrices page and select DD Boost for OpenStorage Compatibility List.

2 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 wide area network (WAN) for offsite disaster recovery.

Unlike virtual tape libraries (VTLs), EMC Data Domain Boost (DD Boost) software does not require any artifacts of tape or tape emulation. Similar to the Data Domain NFS/CIFS implementation, backup images can be written to, and read from, the Data Domain system simultaneously in the applications' native image formats. Backup images can be expired one at a time, leading to improved storage efficiency.

DD Boost increases performance by distributing the deduplication process to the backup server, simplifies disaster recovery procedures, and provides the foundation for additional integration between backup applications and Data Domain systems.

DD Boost software provides the following major functionality. New and enhanced capabilities are available for Single Node and DD Extended Retention.

- [Distributed Segment Processing](#)
- [Managed File Replication](#)
- [Advanced Load Balancing and Link Failover](#)
- [Heterogeneous Application Host Support](#)
- [Virtual Synthetic Backups](#)

- [Storage Quotas](#)
- [Media Server Access Validation](#)

Distributed Segment Processing

The distributed segment processing functionality of the DD Boost software distributes the deduplication process 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 versus sending all of the data over the LAN. The amount of saved network bandwidth depends on the redundant nature of the data being backed up. In general, the greater the redundancy in the data set, the greater the throughput to the Data Domain system.
- With distributed segment processing the DD Boost Library does not cache any data on the application host in memory or to disk, so there are no persistency requirements on the application host. There is no extra memory footprint, but some extra processing power on the application host is required to determine if the data is present on the Data Domain system. Also, because only one file can be concurrently written to or read from each connection, in Backup Exec the number of connections must be the same as the number of files.
- Allows use of existing 1 GbE infrastructure to achieve higher throughput than is physically possible over 1 GbE links. The network bandwidth requirements are significantly reduced because only unique data is sent over the LAN to the Data Domain systems.
- Failed backups can potentially recover much faster. If a large backup fails in the middle or toward the end, the data already sent to the Data Domain system is not resent when the job is restarted by the application. This results in faster completion of the job on retry. The following data flow for a single-controller Data Domain system with the distributed segment processing option enabled shows how DD Boost Library breaks the data

into segments, checks the Data Domain system for segments already stored on it, and sends only unique segments to the Data Domain system. The deduplication process is distributed between the DD Boost Library and the Data Domain system. The DD Boost Library's data handling is transparent to the application. The DD Boost Library does not cache any data on the application host.

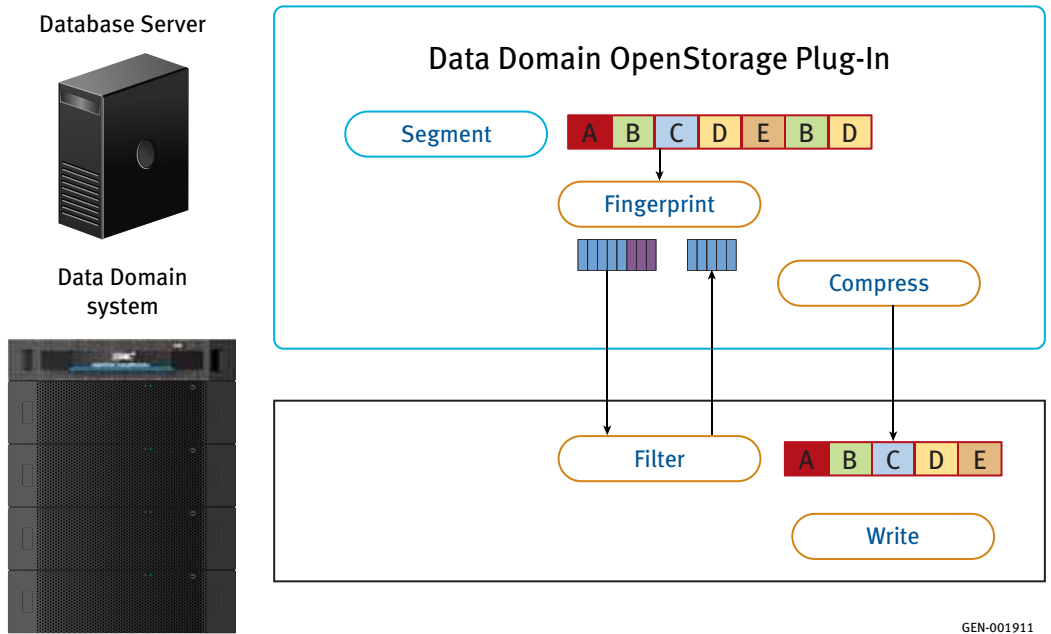


Figure 2-1 Distributed Segment Processing Enabled on a Single Node

The local compression algorithm used by the DD Boost Library is a system-wide configuration setting on the Data Domain system and cannot be configured using the DD Boost Library API. The default local compression algorithm is set to "lz." Data Domain systems support "lz," "gzfast," and "gz" local compression algorithms.

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. Using this approach, the bandwidth required for replication is reduced by up to 99%. This dramatically reduces the time and WAN bandwidth needed to create multiple copies of backups for disk consolidation or disaster recovery purposes. Resource usage on the application servers is also reduced as they are no longer in the data path when duplicate copies of the backup are being created.

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. This design allows backup administrators to create disaster recovery copies of their backups over the WAN using Data Domain Replicator software and keep track of all the copies in the backup application's catalog for easy disaster recovery. This also enables applications to manage retention periods for various copies individually, resulting in a more flexible disaster-recovery deployment.

Leveraging network-efficient Data Domain replication, DD Boost eliminates the need for tape-based backup and recovery procedures at remote sites. Recovery from backup copies at the central site also becomes simplified because all copies are tracked in the backup applications' catalog. Since tape is no longer required at the remote sites, it is entirely practical to deploy Data Domain systems with DD Boost software as the key foundational elements of a multi-site tape consolidation strategy.

Greater control of replication connection bandwidth can be achieved utilizing the current replication throttling settings. For details on managing replication, refer to the *DD OS 5.2 Administration Guide*.

Note: File replication obeys the general throttling rules for replication. While replication is throttled or disabled, the file copy APIs hang until replication is enabled or throttling is removed.

Care must be taken while enabling replication throttling and using managed file replication from clients.

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 5.1. For example, 5.1.0.0, 5.1.0.2, 5.1.1.0, and 5.1.2.0 are all backward compatible.
- Replication is backward compatible across two consecutive release families, such as 5.0 and 4.9, 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.

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 6 Mb/s of available bandwidth.

The low-bandwidth optimization option is available to Data Domain systems with an installed Replicator license. The option is enabled on each Data Domain system and applies to all optimized duplication jobs on that system. Both the source and destination Data Domain systems must be configured with this setting to enable low-bandwidth optimization.

Enabling this option on a Data Domain system is transparent to the backup application. When the backup application requests a Data Domain system to perform a managed file replication job, the source and destination systems automatically perform the

additional compression without involving the requesting backup application.

For more information about this topic, refer to the *DD OS Administration Guide*.

Encrypted Managed File Replication

This option allows applications to encrypt the replication session between two Data Domain systems without requiring additional VPN configuration for securing the WAN connection. When a replication job is initiated between two Data Domain systems by the application using the managed file replication API, encrypted managed file replication uses SSL to encrypt the session between the source and destination Data Domain systems. All data and metadata is sent encrypted over the WAN.

Enabling this option on Data Domain system is transparent to a backup application. When the backup application requests the Data Domain system to perform a replication job, 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.

Encrypted managed file replication is available to Data Domain systems with an installed Replicator license. 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.

Notes:

- For more information about this topic, see the *DD OS 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.

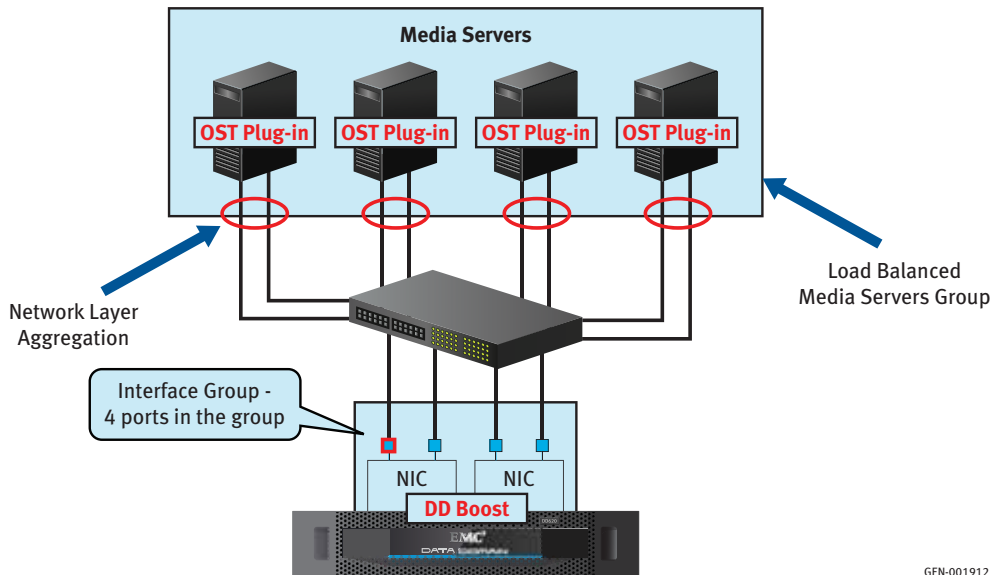
- The low-bandwidth optimization option and the encryption option can be used together.

Advanced Load Balancing and Link Failover

Advanced load balancing and link failover allows combining multiple Ethernet links into a group and registering only one interface on the Data Domain system with the backup application. The DD Boost Library negotiates with the Data Domain system on the interface registered with the application to obtain the best interface to send data to the Data Domain system. Load balancing provides higher physical throughput to the Data Domain system compared to configuring the interfaces into a virtual interface using Ethernet-level aggregation.

The Data Domain system load balances the connections coming in from multiple backup application hosts on all interfaces in the group. Load balancing is transparent to the backup application and is handled by the DD Boost software. Because advanced load balancing and link failover 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 load balancing. 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.



GEN-001912

Figure 2-2 Advanced Load Balancing and Link Failover

Advanced load balancing and link failover can be used with other network layer aggregation and failover technologies. The sample configuration in [Figure 2-2](#) shows one possibility. The links connecting the media servers and the switch connecting the Data Domain system are configured for an aggregated failover mode. This configuration provides network failover functionality from end-to-end. Any of the available aggregation technologies can be used between the media server and the switch.

Advanced load balancing and link failover 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 *DD OS Administration Guide* for more information about how to configure VLAN tagging and IP aliasing on a Data Domain system.

Advanced load balancing and link failover provides the following benefits:

- Eliminates the need to register the Data Domain system on multiple interfaces with the application, which simplifies installation and configuration.
- Prior to DD OS 5.0, any in-process backup job handled by the interface that failed became a failed process that required a restart.
- Starting with DD OS 5.0, all in-process jobs to the failed interface are failed over transparently to healthy operational links. From the point of view of the backup application, the jobs continue uninterrupted.
- If one of the interfaces in the group goes down while the Data Domain system is still operational, the subsequent incoming backup jobs are routed to the available interfaces.
- Backup and restore jobs are automatically load-balanced on multiple interfaces in the group, resulting in higher utilization of the links.
- The Advanced Load Balancing and Link Failover feature must be configured on Data Domain systems. The DD Boost Library APIs do not provide an ability to manage the functionality.
- The interfaces can be added to the group using an IP address only. See the *DD OS Administrator's Guide* for instructions.
- Advanced load balancing and link failover 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 is supported.

In DD OS release 5.2 with DD Boost plug-in 2.5 or higher, advanced load balancing on the Data Domain system network interfaces and failover from one failed network interface to an active network interface has been significantly improved.

In this release an administrator can define multiple network interface groups (ifgroups) where load balancing and failover apply within an ifgroup <group-name>. This increases the capability to support a media server that can reach only some of the Data Domain system interfaces.

This enhancement adds multiple groups to the single default ifgroup supported in previous releases. Load balancing and failover are implemented on a per ifgroup basis, except for the administrative IP. This enhancement is available for Single Node and DD Extended Retention systems.

Each ifgroup <group-name> is composed of a list of interfaces and clients that belong to the ifgroup. Within an ifgroup <group-name> all interfaces are reachable by all the clients for <group-name>.

Each interface or client name may reside in only one ifgroup <group-name>. However, by creating alias IPs the administrator can configure an interface with one IP address in one <group-name> and another IP address for the same interface in another <group-name>. Physical link speed does not affect interface assignments. An ifgroup may consist of mixed 10G and 1G interfaces.

When ifgroup <group-name> is disabled, use the administered IP on the client for all connections, as was done in previous releases. GDA uses management IP for RSS connections.

Only backup and restore jobs reconnect on ifgroup IP for load balancing. All other jobs remain on the administered IP.

New CLI commands have been added under the `ddboost ifgroup` commands to support creation of groups.

Interfaces

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

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

Clients

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

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

Heterogeneous Application Host Support

The combination of a Data Domain system and a backup application that supports DD Boost software creates an optimized connection to provide a tightly integrated solution. DD Boost offers operational simplicity by enabling backup applications to manage optimized connections between the backup application and one or more Data Domain systems.

DD Boost software simplifies the management and allows easy sharing and simultaneous use of disk storage among multiple heterogeneous backup application hosts. This facilitates load balancing and the use of the “best” available backup servers for completing a particular backup job based on pre-configured policies.

DD Boost software supports multiple backup server operating systems providing a seamless experience to both application vendors and application administrators in heterogeneous environments. Further, DD Boost serves as a consolidated interface to Data Domain storage systems, making the systems transparent to the backup application. This simplifies the management of Data Domain deduplication storage systems.

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 DDR. 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 DDR 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 DDR 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 change daily, 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 1 second after the latest incremental. It does NOT include any changes to the backup selection since the latest incremental.

Storage Quotas

Beginning with the 5.2 release of DD OS, users can now use "quotas" to provision Data Domain system storage limits, ensuring that dedicated portions of the Data Domain system are

available as unique storage units. DD Boost storage-unit quota limits may be set or removed dynamically. Quotas may also be used to provision various DD Boost storage units with different logical sizes, enabling an administrative user to monitor the usage of a particular storage unit over time.

Please refer to the `ddbboost`, `quota`, and `mtree` sections of the *EMC Data Domain Operating System 5.2 Command Reference Guide* for details on the quota feature, and commands pertaining to quota operations.

Media Server Access Validation

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

3 Installing DD Boost for OpenStorage

This chapter covers the following topics:

- “Installation Overview” on page 27
- “NetBackup Installation” on page 30
- “Backup Exec Installation” on page 35
- “Tuning Windows Media Servers for Performance” on page 38
- “Uninstalling the Windows Plug-in” on page 38

Note: Complete descriptions of commands used in this guide are provided in the *DD OS 5.2 Command Reference Guide*.

Installation Overview

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

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 optimized duplication. 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, configuration includes specifying the DD Boost user name and password, and creating storage units.

3. Install the OST plug-in software on each media server.
4. After completing the installation steps described in this chapter, configure DD Boost as described in Chapter 4.

OST Plug-In and DD OS Upgrades

The OST plug-in and the DD OS maintain compatibility as defined in the *DD Boost for OpenStorage Compatibility List*. Consult that list before upgrading either the OST plug-in or the DD OS.

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 take advantage of the new functionality available in DD OS. Perform the upgrade as described in [“Installing the OST Plug-In on Media Servers” on page 31](#).

Note: This document illustrates the DD Boost configuration on Data Domain using commands in DD OS 5.2. If you are using a different version of DD OS with this version of the OST plug-in, see the corresponding DD OS version documentation for CLI details.

Firewalls and Ports

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 (Replication)
- TCP 111 (NFS portmapper)
- TCP 2052 (Mount) -- Required for DD OS 5.1, not required for DD OS 5.2

Enabling DD Boost on a Data Domain System

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

▼ Enable DD Boost

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. Set the DD Boost user by entering:

```
# ddboost set user-name ddboost-user
```

5. Enable DD Boost by entering:

```
# ddboost enable
```

```
DD Boost enabled
```

Notes:

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

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

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

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

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

NetBackup Installation

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 storage server's properties.

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

▼ Add these directory locations to the UNIX PATH

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

▼ 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 for uninstalling the OST plug-in on UNIX systems.

▼ Install the UNIX plug-In

1. Download the latest version of the OST plug-in from the Data Domain Support portal.
2. Save the file on the media server in a location of your choice.
3. Stop the backup application's Remote Manager and Monitor Service (nbrmms) process if it is running by entering:

```
# 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 starting 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
```

▼ Correct UNIX plug-in installation or update failure (AIX Media Servers)

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.so
```

```
cp: /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.

If the install script fails:

1. Execute the following command:

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

2. Enter:

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


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

▼ Install the Windows plug-In

The Windows plug-in installer is `libbstspiDataDomainSetup.exe`. It supports 32- and 64-bit Windows plug-ins.

▼ Prepare for installation

1. Download the latest version of the Windows OST plug-in installer from the Data Domain Support portal.
2. 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 running it displays a message to this effect and exits.

3. Stop any NetBackup services. Follow the instructions given in [“NetBackup Windows Services” on page 34](#) to stop the service.
4. Remove any previous plug-in version using either the Windows Control Panel or `libbstspiDataDomainUninstall.exe`.

▼ Start the installation

1. If the services have been stopped, the installer displays the license agreement. Read the terms and Click *I Agree* to continue.
2. In the Choose Install Location dialog box, the correct destination folder is shown. This should not be changed. Click Install to start the installation.

A progress bar monitors the installation.

3. When the Installation is complete, you can click the Show details button to view the files installed and their location.
4. Restart all services. See [“NetBackup Windows Services” on page 34](#).

5. Tune the Windows media server for performance. See [“Tuning Windows Media Servers for Performance”](#) on page 38.

NetBackup Services

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

NetBackup UNIX Services

▼ Stop UNIX services

Enter:

```
# nbrmms -terminate
```

▼ Start or restart UNIX services

Enter:

```
# nbrmms
```

NetBackup Windows Services

▼ Start, stop, or restart Windows services

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.

Backup Exec Installation

This section covers the following topics:

- “Installing the Plug-In on Media Servers” on page 35
- “Backup Exec Services” on page 37

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.

▼ Install the Windows plug-In

The Windows plug-in installer is `libstspiDataDomainSetup.exe`. It supports 32- and 64-bit Windows plug-ins.

▼ Prepare for installation

1. Download the latest version of the Windows OST plug-in installer from the Data Domain Support web site.
2. 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 a service running it displays a message to this effect and exits.
3. Stop any Backup Exec services. Follow the instructions given in “Backup Exec Services” on page 37 to stop the service.
4. Remove any previous plug-in version using either the Windows Control Panel or `libstspiDataDomainUninstall.exe`.

▼ Start the installation

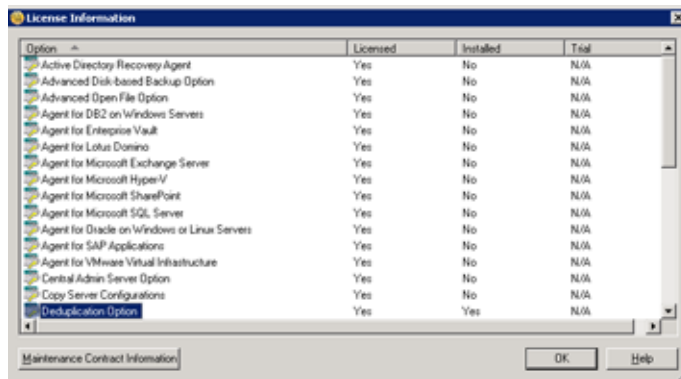
1. If the services have been stopped, the installer displays the license agreement. Read the terms and Click *I Agree* to continue.
2. In the Choose Install Location dialog box, the correct destination folder is shown. This should not be changed. Click Install to start the installation.

A progress bar monitors the installation.

3. When the Installation is complete, you can click the Show details button to view the files installed and their location.
4. Restart all services. See [“Backup Exec Services”](#) on page 37.
5. Tune the Windows media server for DD Boost performance. See [“Tuning Windows Media Servers for Performance”](#) on page 38.

▼ Verify Backup Exec deduplication option

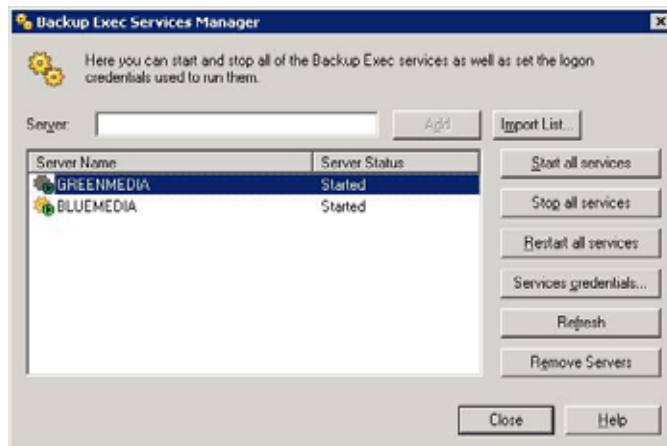
1. From the Backup Exec menu bar, select Help from the About menu.
2. Select License Information. The Deduplication Option should be present in the dialog box. If it is not, re-run the Backup Exec installation and add the Backup Exec Server option.



Backup Exec Services

▼ Start, Stop, or Restart Windows services

1. Within Backup Exec, go to Tools > Backup Exec Services.
2. In the Backup Exec Services Manager window, select the server in the list and click the appropriate services button on the right side of the dialog box.



3. When the processing of the services is completed, click OK in the Backup Exec Services dialog box.

▼ To upgrade the plug-in

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, available on the Data Domain Support portal.

Uninstalling the Windows Plug-in

This procedure applies to NetBackup and Backup Exec.

▼ Uninstall the Windows plug-in

1. Stop the backup application's services.
2. Do one of the following:
 - Uninstall the Data Domain OpenStorage plug-in via 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.

4 DD Boost for OpenStorage Configuration

This chapter covers the following topics:

- “Configuring the Data Domain System” on page 39
 - “Creating Storage Units” on page 40
 - “Configuring Quotas for Storage Units (Optional)” on page 40
 - “Configuring Distributed Segment Processing” on page 41
 - “Configuring Advanced Load Balancing and Link Failover” on page 42
 - “Enabling Low-Bandwidth Optimization ” on page 46
 - “Enabling Encrypted Optimized Duplication” on page 47
 - “Configuring Media Server Access Validation” on page 48
- “Configuring a Media Server” on page 49
 - “NetBackup Configuration” on page 49
 - “Backup Exec Configuration” on page 71

Note: Complete descriptions of commands used in this guide are provided in the *DD OS 5.2 Command Reference Guide*.

Configuring the Data Domain System

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.

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

Creating Storage Units

▼ Create storage units

1. On the Data Domain system, enter:

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

A storage unit name must be unique on any given Data Domain system. However, the same storage unit name can be used on each Data Domain system.

See the section on quotas in the *DD OS 5.2 Command Reference Guide* for details on new command options.

2. Repeat the above step for each Boost-enabled Data Domain system.

Configuring Quotas for Storage Units (Optional)

The storage on a Data Domain system can be provisioned through optional quota limits for a storage-unit. Quota limits can be specified either at the time of creation of a storage-unit, or later through separate commands. For more information refer to the sections on quotas and ddbboost in the *DD OS 5.2 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.

▼ Configure Quotas for Storage Units

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

```
# quota enable
```


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

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

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

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

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

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

Configuring Distributed Segment Processing

The distributed segment processing option is configured on the Data Domain system and applies to all the 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.

Configuring Advanced Load Balancing and Link Failover

If an interface group is configured, when the Data Domain system receives data from the media server clients, the data transfer is load balanced and distributed as separate jobs on the private network, providing higher input/output throughput, especially for customers who use multiple 1 GbE connections. See [“Advanced Load Balancing and Link Failover” on page 19](#) for details.

▼ Configure Advanced Load Balancing and Link Failover

Create an interface group on the Data Domain system by adding existing interfaces to the group and registering the Data Domain system with the backup application, as described below.

1. Create the interface group:

```
# ddboost ifgroup create group_name
```

Examples:

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

Note: The group-name “default” can be used without being created first. In all the remaining `ddboost ifgroup` commands, the “default” group is used if none is specified.

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

```
# ddboost ifgroup external add interface ipaddr
client host-name
```

Examples:

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

# ddboost ifgroup add interface 10.6.109.141
client *
```

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 have one of the interfaces in the ifgroup registered 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 state of the ifgroup, enabled or disabled:

```
# ddboost ifgroup enable status group-name  
  
Status of ifgroup "default" is "enabled"
```

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

```
# ddboost ifgroup show config all
```

Sample output is displayed in the following table.

Group Name	Status	Interfaces	Clients
		Count	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	

After an interface group is set up, you can add or delete interfaces from the group. See [“Modifying an Interface Group” on page 91](#).

To disable a group, force all clients on the group to stop using the ifgroup interfaces with the command:

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

Example:

```
# ddboost ifgroup external disable
```

This action will force all clients *.datadomain.com to use non-ifgroup interfaces.

To move clients *.datadomain to another group, first delete them from the current group then add them to a new group.

Delete clients *.datadomain with the command:

```
# ddboost ifgroup delete group-name interface ip-addr client host
```

Example:

```
# ddboost ifgroup delete external client  
*.datadomain.com
```

Add clients *.datadomain to a new group with the command:

```
# ddboost ifgroup add group-name interface ip-addr  
client host
```

To remove all configuration for a group:

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

Example:

```
# ddboost ifgroup reset external
```

This will remove all clients and interfaces and leave the group disabled.

When 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 host name 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 Advanced Load Balancing and Link Failover (an interface group) from the Enterprise Manager Data Management DD Boost view. (Please see the *DD OS 5.2 Administration Guide*).

Enabling Low-Bandwidth Optimization

To enable the low-bandwidth option, enter:

```
# ddboost file-replication option set low-bw-optim
enabled
```

```
Low bandwidth optimization enabled for optimized
duplication.
```

Note: Enabling or disabling the low-bandwidth optimization option does not require a restart of the Data Domain file system. After enabling low-bandwidth optimization, however, 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. (Please see the *DD OS 5.2 Administration Guide*.)

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

Notes:

- 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 optimized duplication is being done over low-bandwidth networks with less than 6 Mbps aggregate bandwidth.
- For more information on this feature, see the *DD OS 5.2 Administration Guide*.
- The low-bandwidth option for optimized duplication is supported only for standalone Data Domain systems.

Enabling Encrypted Optimized Duplication

To enable the encrypted optimized duplication (encrypted file replication) option, enter:

```
# ddboost 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.

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

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

Directory replication and managed file replication both use the **listen-port** option. Directory replication uses it ONLY to specify the replication destination server listen-port. On the replication source the connection port for a specific destination is entered using the `replication modify` command.

Managed file replication uses replication option **set listen-port** on BOTH the source and destination to specify the port on which the destination listens and the port on which the source connects.

Changing the Replication TCP port will disrupt activity on the Data Domain system. Therefore it should be a planned event.

Notes:

- Enabling or disabling the encrypted optimized duplication option does not require a restart of the Data Domain file system.
- Changing the managed file replication TCP port requires a restart of the Data Domain file system.
- For more information on these topics, see the *DD OS 5.2 Administration Guide*.

Configuring Media Server Access Validation

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

To restrict access, remove the "*" wildcard from the list and add individual clients.

The media 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 (wildcard setting) from the DD Boost access list, enter:

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

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

```
# ddbboost access reset
```

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

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

Example:

```
# ddbboost access add clients ddbboost-dl.emc.com  
ddbboost-dl
```

```
ddbboost-dl.emc.com : Added
```

```
ddbboost-dl : Added
```

To view the DD Boost client access list, enter:

```
# ddbboost access show
```

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

- Wild card "*" followed by partial (for example, "*.emc.com" followed by "*.com")

- Perfect match of sent client name (for example, "ddboost-dl.emc.com")

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

Configuring a Media Server

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

- “NetBackup Configuration” on page 49
- “Backup Exec Configuration” on page 71

NetBackup Configuration

Note: The examples in this chapter assume the following configuration:

- A media server with the name `load64` that runs NBU 6.5 or 7.0.
- Two Data Domain systems with DD Boost enabled named `dd22` and `dd23`.

Media server configuration consists of the following procedures:

- Registering each Data Domain system
- 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

To avoid the need for frequent DNS lookups, add every media server host/IP address to the Data Domain system using the command:

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

DNS reverse lookup has been added for every media server accessing the Data Domain system as of DD OS 4.6.3.

To avoid DNS lookup for every job, also add the Data Domain system IP address into the media server's `/etc/hosts` file.

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 NBU where to direct the operation that queries the plug-in about the storage server's properties.

Concurrent Connection Limit

With NetBackup and Backup Exec, the default number of concurrent connections (jobs) from a single media server is 64. To increase this number, contact your EMC technical consultant.

Registering Each Data Domain System

▼ Register the Data Domain systems

1. On the media server, start the backup application's services. See ["NetBackup Services" on page 34](#).
2. Verify that the plug-in is detected by the backup application by entering:

```
# bpstsinfo -pi
```

The output shows:

- the vendor version, which is the plug-in version
- the build version, which is the OST plug-in version.

3. Before registering a Data Domain system, enable virtual synthetics if that feature is planned to be used, by entering:

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

4. Register a Data Domain system by entering:

```
# 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`.

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

Adding Credentials

▼ Add credentials

1. On a media server that is 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.

2. Repeat the above step for each media server that is to communicate with a specific Data Domain system.

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.

▼ Create a disk pool using the CLI

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
```

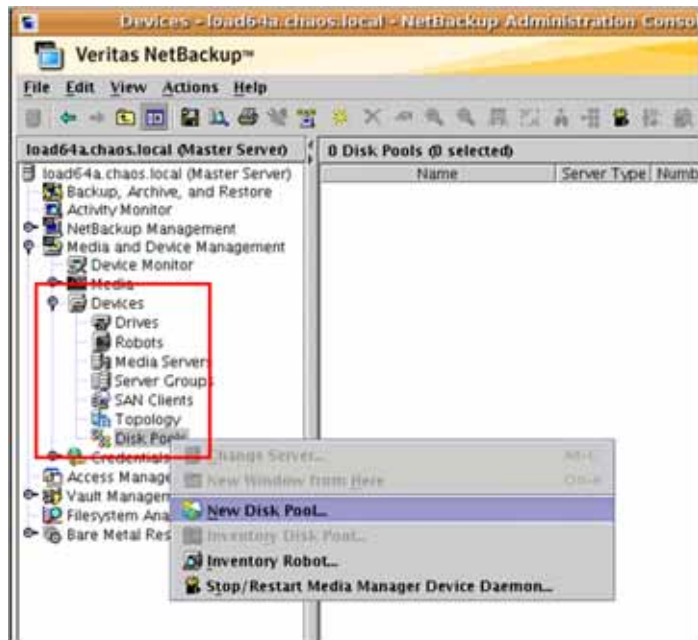
The disk pool name must be unique.

Output similar to the following is displayed:

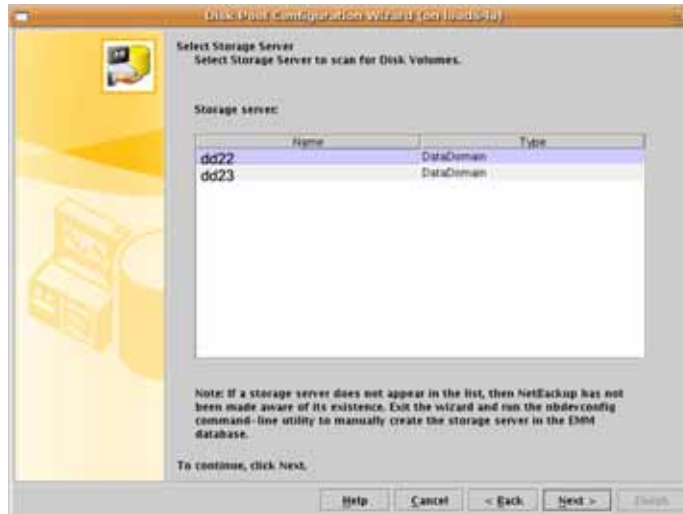
```
Disk pool dd22_storage-unit1_dp has been  
successfully created with 1 volume.
```

▼ Create a disk pool using the GUI

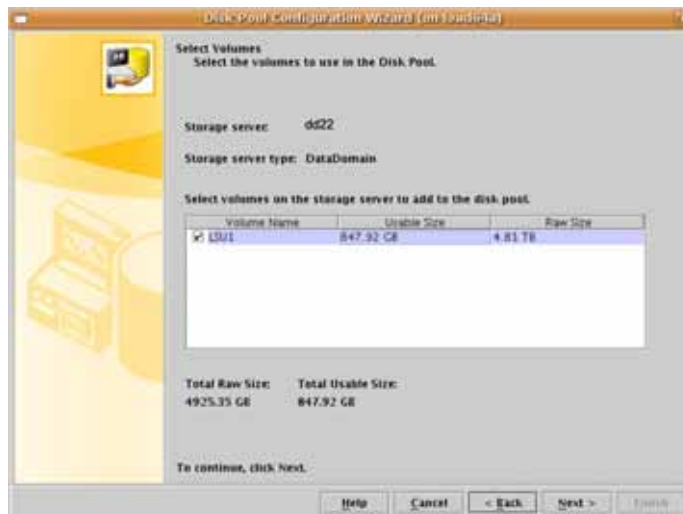
1. Open the NetBackup Administration Console's Devices window.



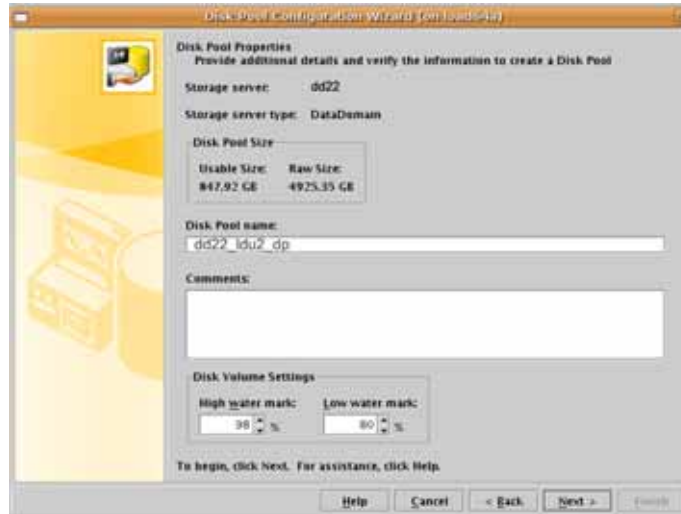
2. In left pane of this window, expand the Media and Device Management node. Expand the Devices node. Right-click Disk Pools. Select New Disk Pool from the menu.
3. The Disk Pool Configuration Wizard displays. Select the Data Domain system from the Storage server list, and click Next.



4. Select (check) the storage unit in the Volume Name list, and click Next.



5. In the Disk Pool Properties window, enter a unique name for the disk pool in the Disk Pool name text box. Click Next.



6. After creating the disk pool, you are given the option of continuing to use the GUI to create a storage unit, which is a collection of disk pools. For instructions, see [“Creating Storage Units”](#) on page 55.

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.

▼ Create storage units using the CLI

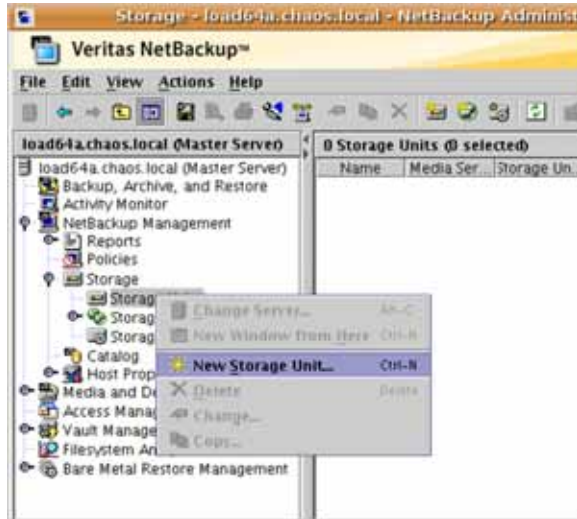
Enter a command similar to the following:

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

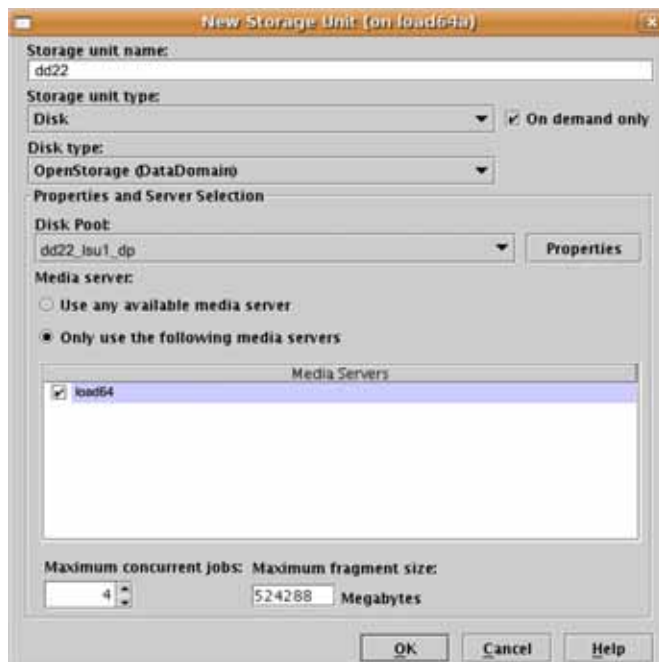
Note: There is no output from this command.

▼ Create storage units using the GUI

1. Open the NetBackup Administration Console's Storage window.



2. In the left pane, expand the NetBackup Management node. Under the Storage node, right-click Storage Units, and select New Storage Unit from the menu.
3. In the New Storage Unit dialog box, enter a unique name for the new storage unit.



4. For the Storage unit type, select Disk from the list.
5. For the Disk type, select OpenStorage (DataDomain) from the list.
6. For the Disk Pool, select the desired disk pool from the list.
7. Select the media servers that can access this storage unit.
 - If any media server is allowed to access the storage unit, select Use any available media server.
 - If a storage unit needs to be visible to a select set of media servers, select (check) the name of each in the Media Servers list.

Note: If a server does not appear in the list, verify that its OST plug-in has been installed and that its login credentials have been configured.

8. Click OK.

Creating a Backup Policy

For instructions on creating a backup policy, see the *Symantec NetBackup Administration Guide*.

Configuring Buffers

You can set the number and size of various buffers, but 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 the `SIZE_DATA_BUFFERS` and `SIZE_DATA_BUFFERS_DISK` to 262144.

▼ Create buffer files

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.

Notes:

- Prior to NetBackup version 6.5.4, when an optimized duplication job failed, duplication job retry attempted to use conventional duplication. This resulted in sending the fully inflated complete backup file from the source OpenStorage storage unit through one, or possibly two, NetBackup media servers to the destination OpenStorage storage unit.
- NetBackup version 6.5.4, by default, does not retry a failed optimized duplication job unless it was initiated by means of a Storage Lifecycle Policy. When retried via a Storage Lifecycle Policy, the job attempts to use optimized duplication only. The default behavior can be overridden by configuring NetBackup to retry with conventional duplication. See the *Veritas NetBackup™ Shared Storage Guide UNIX, Windows, Linux Release 6.5.4* document for detailed information.

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.

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.

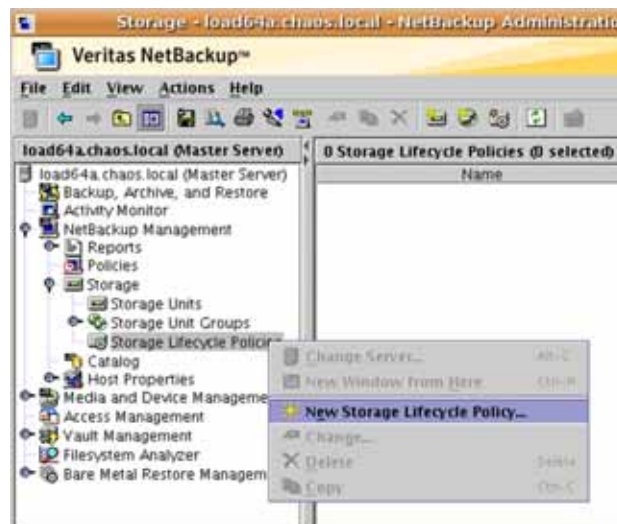
Notes:

- 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.

▼ Create a storage lifecycle policy

This example creates a lifecycle policy that backs up to a DD Boost-enabled Data Domain system and then creates a duplicate copy on another DD Boost-enabled Data Domain system. There are two storage destinations. The first serves as backup to the first DD Boost-enabled Data Domain system. The second storage destination is used for duplication to the second DD Boost-enabled Data Domain system.

1. Open the NetBackup Administration Console's Storage window.



2. In left pane of the window, expand the NetBackup Management node and then the Storage node. Right-click

Storage Lifecycle Policies. Select New Storage Lifecycle Policy from the menu.

3. In the New Storage Lifecycle Policy dialog box, enter a lifecycle policy name in the Storage lifecycle policy name text box.

New Storage Lifecycle Policy (on loadbda)

Storage lifecycle policy name:

Data classification:

Duplication job priority (higher number is greater priority):

Use for	Storage Unit	Volume Pool	Media Owner	Retention Type	Retention Period	Alternate Read S...
Backup	dd22_lsu1_su	--	--	Fixed	1 week	--
Duplication	dd23_lsu1_su	--	--	Fixed	1 week	--

Buttons: Add... Change Remove OK Cancel Help

4. Click the Add button to add a new storage destination to the list.
5. In the New Destination dialog box, select either Backup or Duplication, depending on how the policy is to be used. In this example, select Backup.

New Destination (on loadbda)

Use for: ☒ Backup ☐ Duplication

Retention Type: ☒ Fixed ☐ Staged capacity managed ☐ Expire after duplication

Retention period:

Desired cache period:

Storage unit:

Volume pool:

Media Owner:

Alternate read server:

Buttons: OK Cancel Help

6. Select the storage unit (dd22_storage-unit1_su) from the list.

7. Select one of the retention types:

- Fixed

Fixed retention means that the backup is retained for a specific period of time. Afterwards, NetBackup deletes information about the expired backup, and the files in the backup become unavailable to restore operations. Fixed retention enables you to use different retention periods for backup copies on local and remote Data Domain systems. For example, you might retain backups at the local Data Domain system for one week and the copies at a remote disaster recovery Data Domain system for four weeks.

- Staged capacity managed

This type means that the Data Domain system serves as a staging device before it is duplicated to a final destination. The Data Domain system must be configured to use disk staging.

- Expire after duplication

This type applies only to backup destinations, not to duplication destinations. As soon as NetBackup duplicates the backup, NetBackup deletes information about the backup and the files in the backup become unavailable to restore operation.

8. Click OK in the New Destination dialog box.

9. Return to the New Storage Lifecycle Policy dialog box.

10. Click the Add button.

11. In the New Destination dialog box, select a different storage unit for Duplication. Click OK, which adds it to the list in the New Storage Lifecycle Policy dialog box.

12. Click OK.

Note: NetBackup 6.5.4 introduced Hierarchical Duplication, which allows creation of up to ten duplicate copies using Storage Lifecycle Policies. Policies can be defined such that the third copy can be sourced from the second copy, which is the first duplicate copy. See the NetBackup documentation for more information.

Configuring a Virtual Synthetic Backup

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

1. Virtual Synthetics is disabled by default. Enable a virtual synthetic backup on the Data Domain system by entering:

```
# ddboost option set virtual-synthetics {enable | disable}
```

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
```

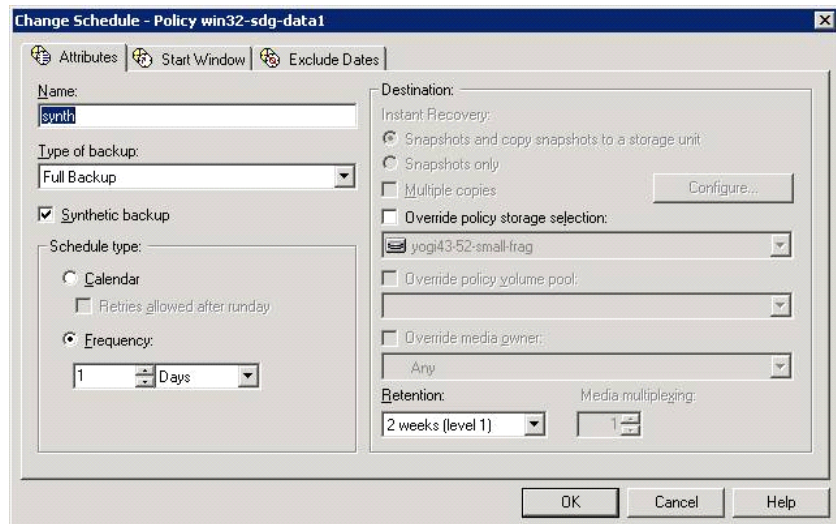
If the `OptimizedImage` flag is not displayed in the output, configure it with the `nbdevconfig` command:

```
# nbdevconfig -changests
```

3. Configure NetBackup policy attributes and schedule details.

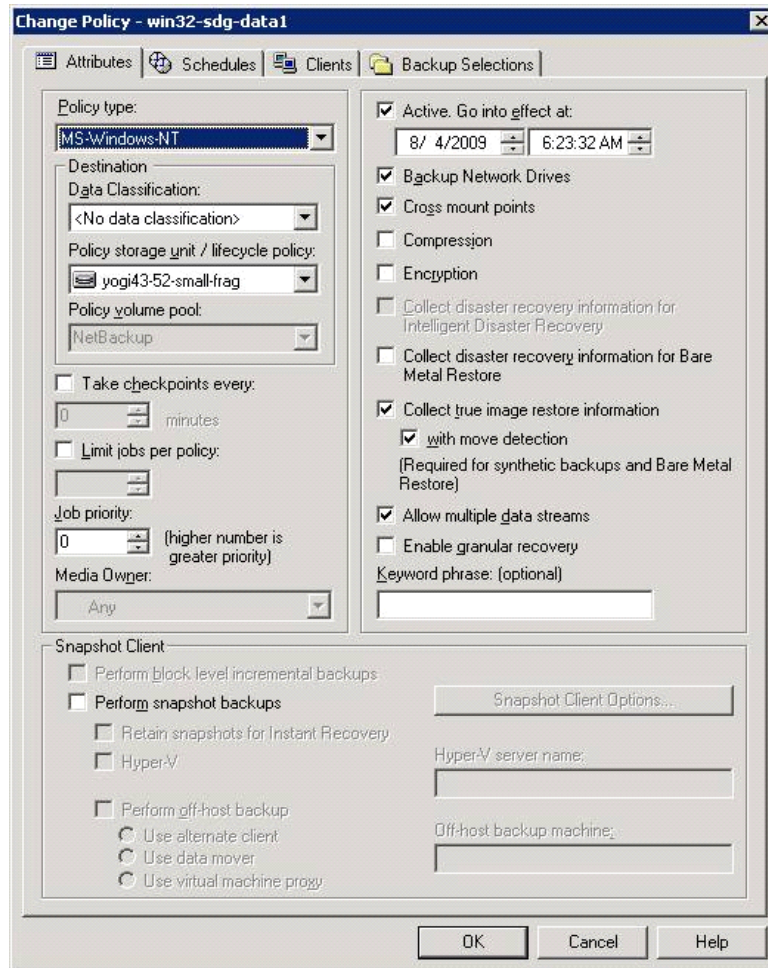
In the Change Schedule window:

- Add a new schedule with a user-specified name.
- Check `Synthetic backup`.



4. In the Change Policy window:

- Add a new schedule with a user specified name.
- Uncheck Bare Metal Restore.
- Check Collect true image restore information
- Check with move detection



Note: 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.

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_CI_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_CI_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_CI_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_CI_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_CI_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_CI_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-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
rtp-ost-sparcl.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-sparcl.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-sparcl.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
-----
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
-----
```

Backup Exec Configuration

When you launch the Backup Exec application, its Home page is displayed. This page guides you through the configuration procedures. It provides links to Backup Exec technical support and documentation, as well as an alerts and job history summaries.

Note: For information on setting up, scheduling, and monitoring jobs, see the *Symantec Backup Exec 2010 Administrator's Guide*.

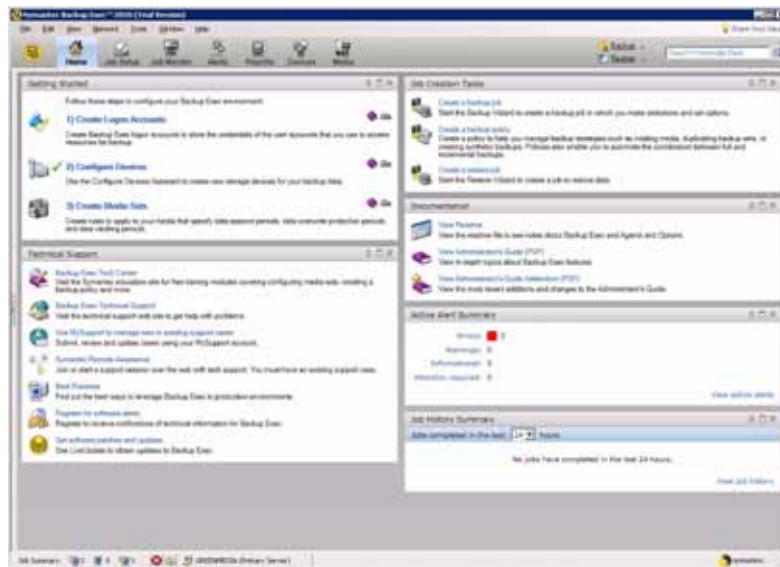


Figure 4-1: Backup Exec Home Page

Complete these Getting Started steps (as shown in the upper-left of the home page):

1. Create a logon account.
2. Configure devices.

Notes:

- Click the icon to the left of the task to launch a wizard or assistant that guides you through the configuration step.
- A green check to the left of a task indicates that a task has been completed.

Creating a Logon Account

Follow these steps to create a logon account.

▼ Create a logon account

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.

▼ Configure devices

1. Create a storage unit on the Data Domain system. See [“Creating Storage Units” on page 40](#).
2. From the Backup Exec Home page, select Configure Devices from the Tools menu. Select Add OpenStorage from the menu.

3. Configure the AddOpenStorage Device dialog box's General tab as follows:
 - Name: Enter the name of the Data Domain system.

The screenshot shows the 'Add OpenStorage Device' dialog box with the 'General' tab selected. The fields are filled with the following values: Name: rtp-ddr8-120, Server: rtp-ddr8-120.datadomain.com, Logon account: sysadmin, Server type: DataDomain, and Logical storage unit: greenmedia. Below the storage unit, a note states: '(Represents a single storage element, such as a volume or directory)'. The 'Concurrent Operations' section has a spin box set to 15, with the text 'Allow 15 concurrent operations for this device.' at the bottom. The dialog box has OK, Cancel, and Help buttons at the bottom right.

- 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.

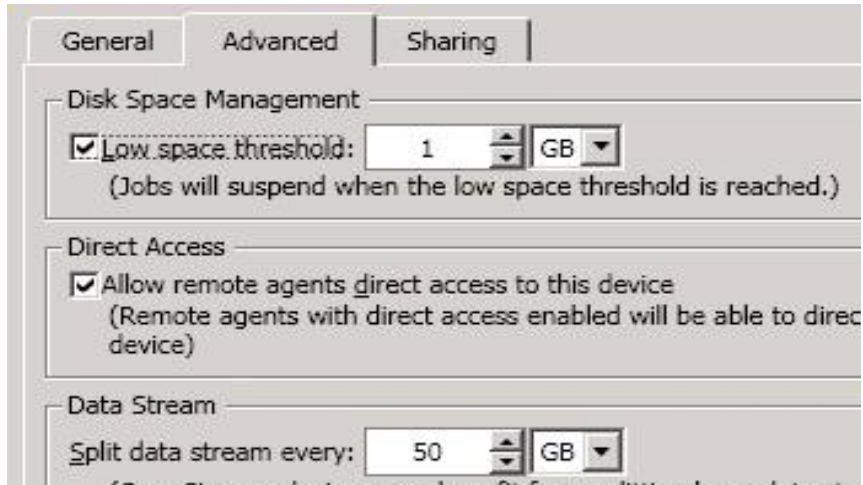
The total number of concurrent connections (jobs) from a single media server OST plug-in to all associated OpenStorage storage units is 48. The concurrent operations limit for a single device can be determined as follows:

$$48 \geq \# \text{ OpenStorage storage units} + \sum \text{ concurrent operations for each storage unit}$$

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 AddOpenStorage 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.

After the device has been configured, the new storage unit is displayed in the Devices page (Figure 4-2).




Figure 4-2: Backup Exec Devices

Configuring Optimized Duplication

The ways to develop duplication jobs in Backup Exec are described in detail in the *Symantec Backup Exec 2010 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.

Data Domain 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.

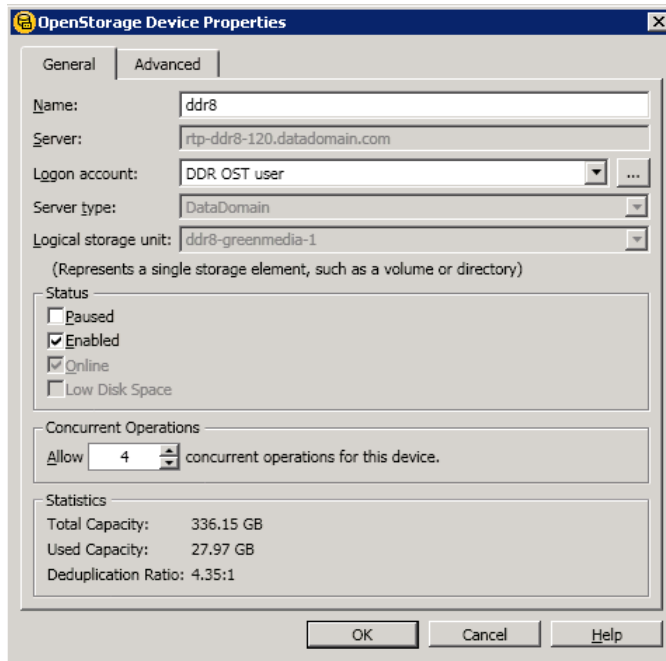


Figure 4-3: OpenStorage Device Properties (General Tab)

This value is set during storage unit creation when the media device representing the storage unit on the Data Domain system is added to Backup Exec.

- The Concurrent Operations count of the destination Data Domain system is greater than or equal to that of the source Data Domain system. See [Figure 4-3](#).

Configuration Limitations for Optimized Duplication

1. Optimized Duplication is supported with Backup Exec 2010 R2 or higher.
2. 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.

Figure 4-4 show an example of the type of backup images that can be duplicated using optimized duplication.

OST-00000129-446E7A3EBD620BBE can be duplicated using optimized duplication, whereas

OST-00000128-446E7A3EBD620BBE cannot be duplicated because it contains data from multiple sources.

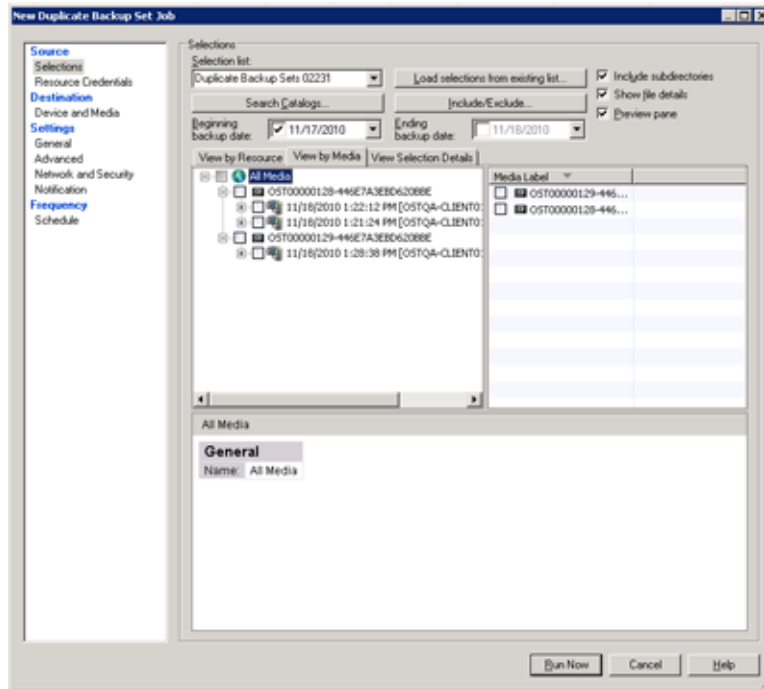
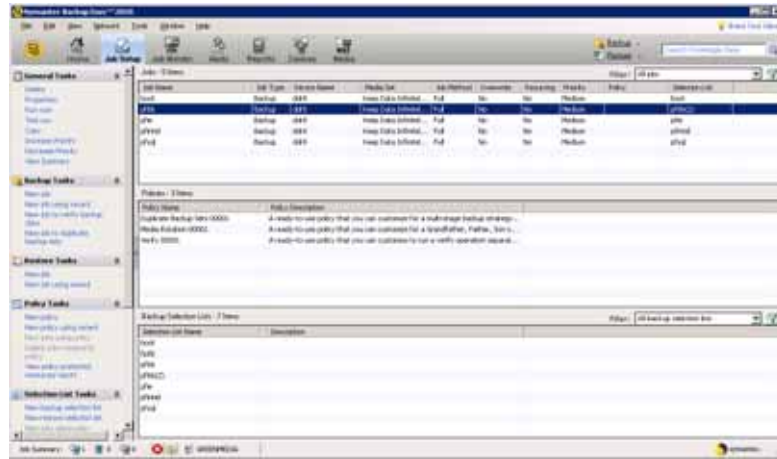


Figure 4-4: New Duplicate Backup Set Job


▼ Add a duplication job to an existing job

These instructions assume that you have developed a backup job that is running (or scheduled to be run) in order to generate a backup set on the source Data Domain system (*media device ddr5*).

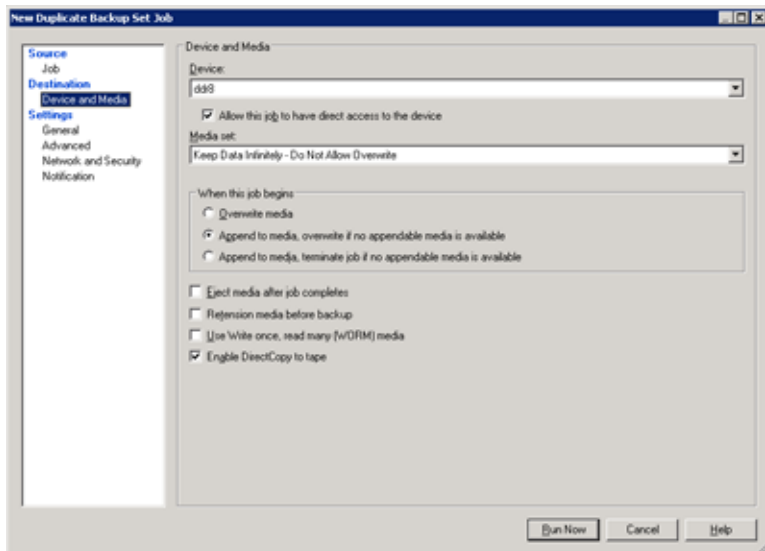
1. Open the Job Setup panel.



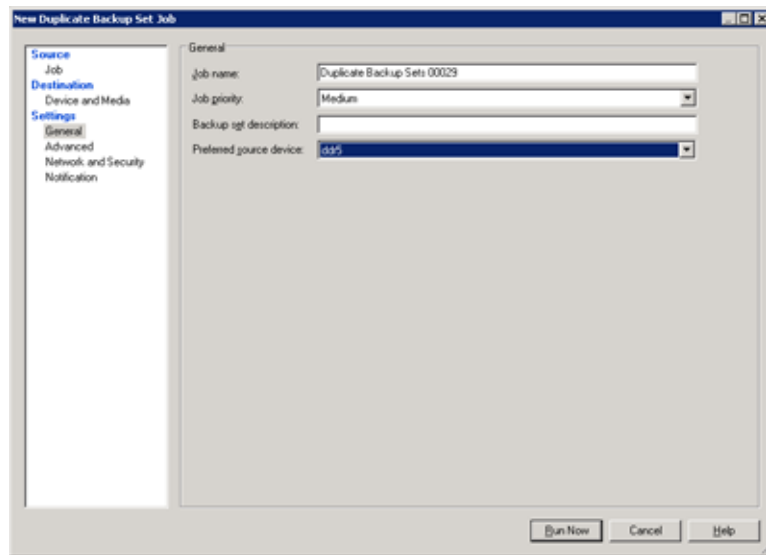
2. In the left panel, select Backup Tasks > New job to duplicate backup sets.
3. In the New Job to Duplicate Backup Sets dialog box, select the *Duplicate backup sets following a job* option, and click OK.
4. In the Source > Job panel, select an available job that is running or scheduled to be run, *pf86*, as shown below.



5. In the Destination > Device and Media panel, select the remote Data Domain system, in this case a device named *ddr8* as shown below. Complete file options as required.




6. In the Settings > General panel, verify that the source Data Domain system is set properly, in this example a device named *ddr5* as shown below.




The Job Monitor panel displays the status of the jobs as they are scheduled and run.

The scheduled backup job, *pf86*, and a scheduled duplicate backup job are shown below:



Each time the backup job runs, the duplicate job is run after it.



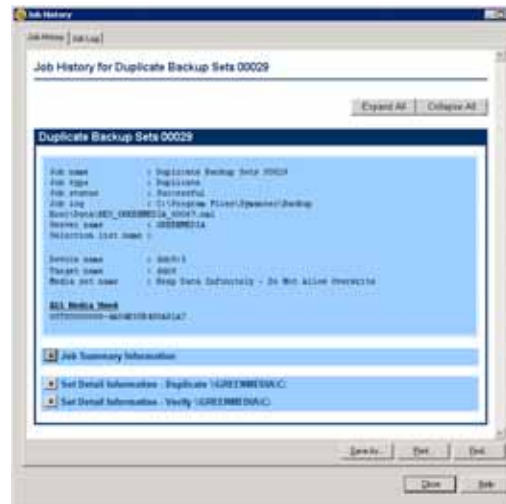



Figure 4-5: Job History

The Job History dialog box shown in [Figure 4-5](#) contains job statistics and log events for each of the two jobs.

▼ Duplicate a previous job

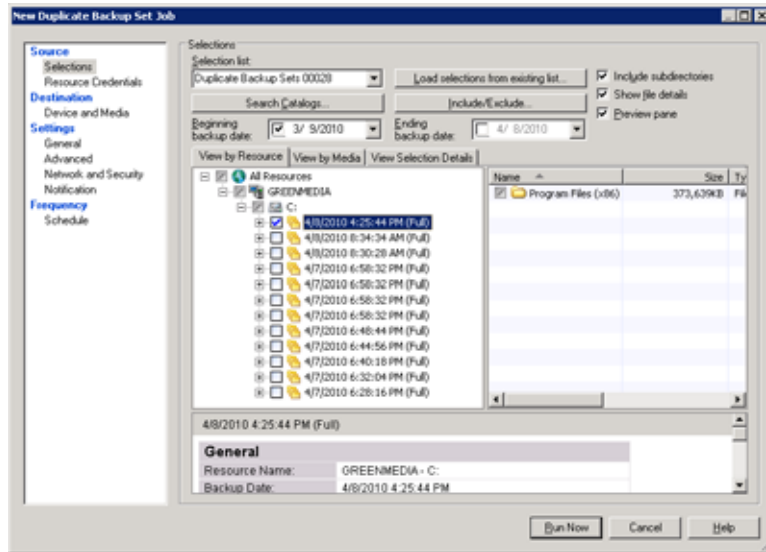
Prerequisite: An existing backup job and at least one backup set saved on the source Data Domain system (media device named *ddr5* in this example).

1. Open the Job Setup panel.

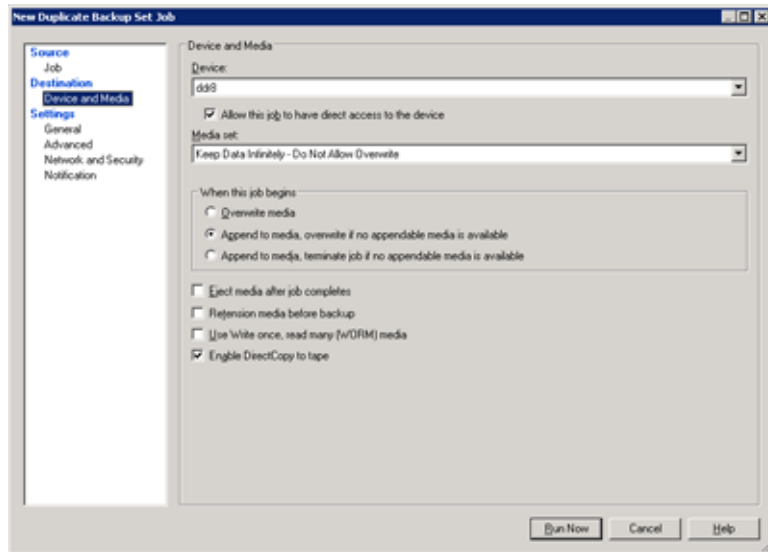


2. Select the New job to duplicate backup tasks in the left hand sub-panel:
3. In the New Job to Duplicate Backup Sets, select the Duplicate existing backup sets option, and click OK.


4. In the Source > Selections panel, select an available media set that has been created on media device *ddr5*.



5. In the Destination Device and Media panel, select the remote Data Domain system, in this case a device named *ddr8*. Complete file options as required.




6. In the Settings > General panel, verify that the source Data Domain system is set properly, in this example a device named *ddr5*.




7. You can schedule this duplication for later, or click Run Now to start it.

The Job Monitor panel displays the status of the job as it runs.



Upon completion the job is moved to the Job History list.



The Job History contains job statistics and log events.

5 Administration

This chapter covers the following major topics:

- “Data Domain System Administration” on page 89
- “NetBackup Administration” on page 94
- “Backup Exec Administration” on page 96

Note: Complete descriptions of commands used in this guide are provided in the *DD OS 5.3 Command Reference Guide*.

Data Domain System Administration

This section covers the following topics:

- “Restricting Export Permissions to DD Boost” on page 89
- “Optimized Duplication Version Compatibility” on page 91
- “Modifying an Interface Group” on page 91
- “Removing Advanced Load Balancing and Link Failover Configuration” on page 93

Restricting Export Permissions to DD Boost

Data Domain systems export the default mount created for DD Boost `/backup/ost` as follows:

```
/backup/ost *  
(rw,no_root_squash,no_all_squash,insecure)
```

By default, with this export definition, the DD Boost service is accessible to all media server clients. To restrict access to a certain set of media servers, remove this default export definition and add a new export definition to specify only those media servers.

This example shows removing the default export and adding a new export definition with access restricted to two specific media servers, `mediaserver1.datadomain.com` and `mediaserver2.datadomain.com`:

```
ddboost disable
nfs del /backup/ost *
nfs add /backup/ost mediaserver1.datadomain.com
mediaserver2.datadomain.com
    (rw,no_root_squash,no_all_squash,insecure)
ddboost enable
```

This set of steps creates an export entry that enables DD Boost access only to the two specified media servers. You must explicitly specify the export options (`rw,no_root_squash,no_all_squash,insecure`) that are to override the default export option of `secure`. The `insecure` option is required to allow the DD Boost client software to use a broader range of host-side TCP port numbers.

If there is no specific export entry established, the `ddboost enable` operation installs the default export entry, which allows all hosts to access the DD Boost service. However, if there is any export entry already established for `/backup/ost`, the `ddboost enable` operation does not modify the export list, thus allowing the already-installed export controls to be honored.

Prior to DD OS 4.9.1, the default `/backup/ost *` entry was always added during `ddboost enable`.

Notes:

- Disabling DD Boost disables data access to all media servers.
- Administrators need to ensure that no backup jobs are running to the Data Domain system during these operations.
- A fully-qualified domain name, an IP address, or a DNS resolvable name of the client must be specified during the modification of the client access list.
- The `rw` option is required on the export entry to allow the media server to write backup data to the Data Domain system.

Optimized Duplication Version Compatibility

The Data Domain policy of upgrade compatibility for optimized duplication 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 5.1. For example, 5.1.0.0, 5.1.0.2, 5.1.1.0, and 5.1.2.0 are backward compatible.
- Optimized duplication is backward compatible across two consecutive release families, such as 4.9 and 5.0, although only the current release within each family is tested.
- The destination Data Domain system must be at the highest version. Cases in which the destination node is at a lower revision than the source node or nodes are not supported.
- For bidirectional optimized duplication, both destination and source must run the same release.
- Optimized duplication over WAN should use the Data Domain `replication` command to adjust for low bandwidth and long latency.
- To configure encrypted optimized duplication, both the source and the destination must run DD OS 5.0 or later.

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.

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

```
# ddbboost show connections
```

Sample output is from a single-node Data Domain system. See the DD860 Archiver Administration Guide for an example of this command on an active tier.

ifgroup		Connections							
Group	Name	Status	Interfaces	Bckup	Restore	Src-repl	Dst-repl	Synthetic	Total

none			10.6.109.40	0	0	0	0	0	0
none			10.6.109.41	0	0	0	0	0	0
none			10.6.109.244	0	0	0	0	0	0
none			10.6.109.200	0	0	0	0	0	0
none			10.6.109.235	0	0	0	0	0	0
default	enabled		10.6.109.144	0	0	0	0	0	0
default	enabled		10.6.109.145	0	0	0	0	0	0
10Glab	enabled		10.6.109.230	0	0	0	0	0	0

Total Connections:									

				0	0	0	0	0	0

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

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

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

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.

If you make changes using the `net` command that modify the interfaces, such as enabling an interface that is configured for ifgroup, execute the `ddbboost show connections` command to update the load balancing view. This updating allows the ifgroup to use the interface.

Removing Advanced Load Balancing and Link Failover Configuration

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

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:

```
# ddbboost 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:

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

4. Reset the interface group:

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

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.

NetBackup Administration

This section covers the following topics and procedures:

- “Find your OST plug-in version” on page 94
- “Find your NetBackup version” on page 94
- “Network Time-Outs” on page 94
- “Grouping Storage Units to Provide Failover” on page 95

▼ Find your OST plug-in version

Enter:

```
# bpstsinfo -pi
```

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

▼ Find your NetBackup version

NetBackup

Search the file `[NetbackupInstall_Dir] version`.

Sample output from this command looks like:

```
HARDWARE SOLARIS
VERSION NetBackup? 6.5
RELEASEDATE Mon Jul 23 16:30:00 CDT 2009
BUILDNUMBER 20070723
```

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

Add the following two lines to the file `[NetBackupInstall_directory]/bp.conf`:

```
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

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

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

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

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

2. Delete all of the policies from the GUI.

3. Delete all of the storage units by entering:

```
# bpstudel -label SU_NAME
```

4. Delete all the disk pools by entering:

```
# nbdevconfig -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 100](#).

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

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

Backup Exec Administration

This section covers the following topics:

- [“Find your OST plug-in version” on page 97](#)
- [“Find your Backup Exec version” on page 97](#)
- [“Delete storage units on Data Domain Systems” on page 97](#)
- [“Recover backup files on an OpenStorage device” on page 97](#)

▼ Find your OST plug-in version

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

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

▼ Delete storage units on Data Domain Systems

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-name>
```

▼ Recover backup files on an OpenStorage device

1. From the Backup Exec Home page, select the Devices tab.
2. Select the Data Domain system storage unit on the Devices page.
3. Select a media label.
4. Right-click the media label, and select Restore from the menu.

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.

This chapter covers the following topics:

- “General Troubleshooting” on page 99
- “Data Domain System Settings for File-Replication” on page 100
- “NetBackup Troubleshooting” on page 100
- “Backup Exec Troubleshooting” on page 109

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

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 *DD Boost for OpenStorage Compatibility List*. A supported configuration can become unsupported if any component changes.

This list is available from the Data Domain Support portal (<https://my.datadomain.com>). Navigate to the Documentation > Compatibility List page and select DD Boost Compatibility List.

- 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 49](#) 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 lead to 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:

1. Enter:

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

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

On a Windows System

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 103](#).

On a UNIX System

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 103](#).

▼ Check the installation

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

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
```

a. 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/openv/netbackup/logs/admin` directory.

▼ Check credentials

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
```

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 51](#), 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:

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

Note: An extra entry can occur after upgrading NetBackup from version 6.0 to 6.5.

Error Logging on the Media Servers

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

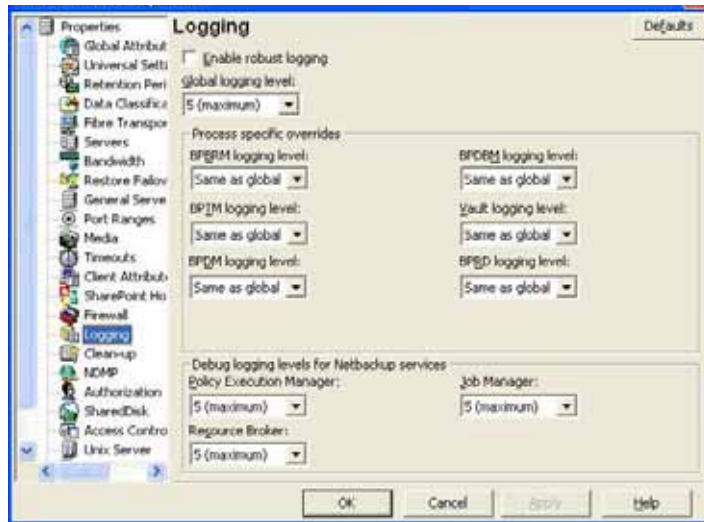
▼ Enable error logging on a Windows system

Enter:

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

▼ Enable logging via the GUI

1. Open the NetBackup Administration Console's Devices window.
1. In left pane of this window, expand the NetBackup Management node.
2. Expand Host Properties and select Master Servers.
3. In the right pane, double-click the machine name.
4. In the Property directory tree, select Logging.



5. Set the Global Logging level to 4, and click OK.

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

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

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

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 file-replication:
`C:\Program Files\Veritas\logs\bpdm`

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

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.

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 58](#).

- Program Not Registered

The following output indicates that the program is not registered:

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

2. Enable DD Boost by installing a valid license:

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

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

```
# fileysys status
```

▼ Resolve Cannot connect on socket error

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

Follow these steps:

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

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

NetBackup Backup Jobs Fail on Solaris Media Servers

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

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

2/28/2009 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> apodddrrp01:
/usr/opensv/lib/ost-
plugins/libstspiDataDomain.so:stspi_open_server
STS_EPLUGIN Can't connect to mountd on apodddrrp01
(: 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

1. Uncomment or remove `sun-dr` entries in

```
/etc/inet/inetd.conf:
```

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

```
sun-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:

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

```
ipseccnf -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 Replicator

1. Obtain a replicator 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-name]`
- Verify that DDP_SynWR RPCs are being sent.
`ddboost show stats`
- Verify that OptimizedImage flag is set.
`nbdevquery -liststats`
- Verify virtual-synthetics is enabled on the Data Domain system.
`ddboost optionj show`

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 16. This number was specified when you configured the device. See [“Configuring Devices” on page 72](#).
- 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.

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

1. From the Backup Exec Home page, go to the Devices page.
2. Right-click the OpenStorage device and display its Properties. Note the logon account.
3. From the system menu go to Network > Logon Account Management.
4. Select the logon account and verify its user name. Try to edit the account to verify the password. This user name and password should match the user name set for DD Boost and password on the Data Domain system.

▼ Resolve license errors

1. If the OpenStorage option is not enabled, from the Backup Exec Tools menu select Install Options and License Keys on this Media server.
2. Click Next to go to Configure Options. See [Figure 6-1](#).
3. Open the Backup Exec Options and verify that the Deduplication Option is selected.
4. Click Cancel to avoid making unintended changes.

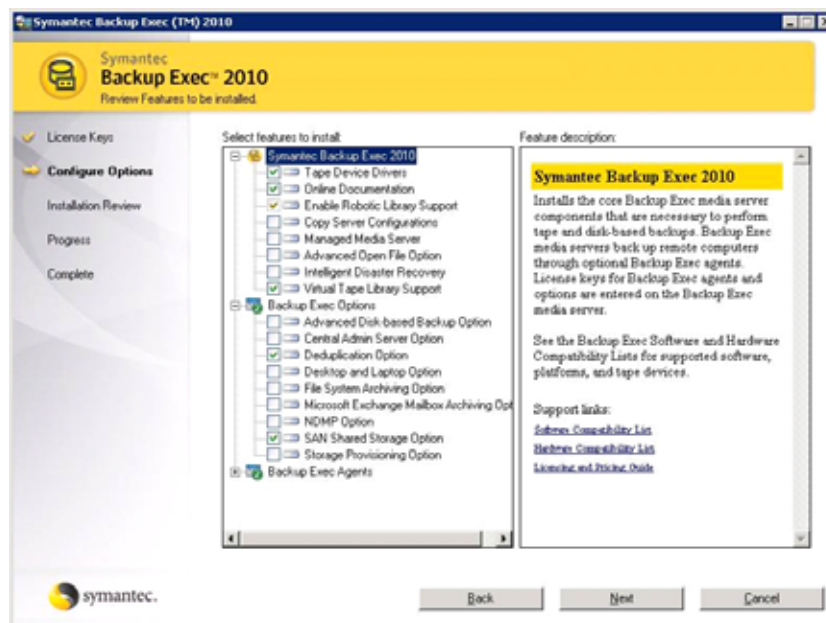


Figure 6-1: Backup Exec Review Features for Installation Page

▼ Set up active debugging

Use the Backup Exec debugging utility to troubleshoot Backup Exec issues.

1. From the Backup Exec Home page, from the Tools menu select Support Utilities > Run the Debug Monitor for Active Debugging.

The Backup Exec Debug Monitor is displayed.




Figure 6-2: Backup Exec Debug Monitor

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. From the Tools > Settings menu, select the Device and Media Debug setting (see [Figure 6-3](#)). Select *Enable verbose logging* if it is not enabled.




Figure 6-3: Settings Menu

The log file that is captured is named *media server-SGMon.txt*. It is located in the logs directory.

