

Symantec NetBackup™ Deduplication Guide

UNIX, Windows, Linux

Release 7.5

Symantec NetBackup™ Deduplication Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version: 7.5

PN: 21220065

Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo, Veritas, and NetBackup are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our Web site at the following URL:

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America supportsolutions@symantec.com

Contents

Technical Support	4	
Chapter 1	Introducing NetBackup deduplication	13
	About NetBackup deduplication	13
	About NetBackup deduplication options	13
	How deduplication works	15
	New features and enhancements for NetBackup 7.5	16
Chapter 2	Planning your deployment	19
	Planning your deduplication deployment	20
	About the deduplication tech note	21
	NetBackup naming conventions	22
	About the deduplication storage destination	22
	About the NetBackup Media Server Deduplication Option	23
	About NetBackup deduplication servers	25
	About deduplication nodes	26
	About deduplication server requirements	26
	About media server deduplication limitations	28
	About NetBackup Client Deduplication	28
	About client deduplication requirements and limitations	30
	About remote office client deduplication	31
	About remote client deduplication data security	31
	About remote client backup scheduling	32
	About NetBackup Deduplication Engine credentials	32
	About the network interface for deduplication	33
	About deduplication port usage	33
	About deduplication compression	34
	About deduplication encryption	34
	About optimized synthetic backups and deduplication	35
	About deduplication and SAN Client	36
	About optimized duplication and replication	36
	About MSDP optimized duplication within the same domain	36
	About NetBackup Auto Image Replication	47
	About deduplication performance	52

- How file size may affect the deduplication rate 54
- About deduplication stream handlers 54
- Deployment best practices 54
 - Use fully qualified domain names 54
 - About scaling deduplication 55
 - Send initial full backups to the storage server 55
 - Increase the number of jobs gradually 56
 - Introduce load balancing servers gradually 56
 - Implement client deduplication gradually 57
 - Use deduplication compression and encryption 57
 - About the optimal number of backup streams 57
 - About storage unit groups for deduplication 57
 - About protecting the deduplicated data 58
 - Save the deduplication storage server configuration 59
 - Plan for disk write caching 59
 - How deduplication restores work 59
- Replacing the PureDisk Deduplication Option with Media Server
 - Deduplication on the same host 60
- Migrating from PureDisk to the NetBackup Media Server
 - Deduplication option 61
- Migrating from another storage type to deduplication 62

Chapter 3 Provisioning the storage 65

- About provisioning the deduplication storage 65
- About deduplication storage requirements 66
 - About support for more than 32-TB of storage 68
- About deduplication storage capacity 68
- About the deduplication storage paths 69
- Do not modify storage directories and files 69
- About adding additional storage 70
- About volume management for NetBackup deduplication 70

Chapter 4 Licensing deduplication 71

- About licensing deduplication 71
- About the deduplication license key 72
- Licensing NetBackup deduplication 72

Chapter 5 Configuring deduplication 75

- Configuring NetBackup media server deduplication 76
- Configuring NetBackup client-side deduplication 78
- Configuring a NetBackup deduplication storage server 79

About NetBackup deduplication pools	79
Configuring a deduplication disk pool	81
Media server deduplication pool properties	81
Configuring a deduplication storage unit	83
Deduplication storage unit properties	85
Deduplication storage unit recommendations	86
Enabling deduplication encryption	87
Configuring optimized synthetic backups for deduplication	89
About configuring optimized duplication and replication bandwidth	89
Configuring MSDP optimized duplication copy behavior	91
Configuring a separate network path for MSDP optimized duplication	92
Configuring optimized duplication of deduplicated data	94
About the replication topology for Auto Image Replication	97
Configuring a target for MSDP replication	98
Viewing the replication topology for Auto Image Replication	100
Sample volume properties output for MSDP replication	101
About the storage lifecycle policies required for Auto Image Replication	102
Customizing how nbstserv runs duplication and import jobs	105
Creating a storage lifecycle policy	105
Storage Lifecycle Policy dialog box settings	106
Adding a storage operation to a storage lifecycle policy	108
About backup policy configuration	110
Creating a policy using the Policy Configuration Wizard	111
Creating a policy without using the Policy Configuration Wizard	111
Enabling client-side deduplication	112
Resilient Network properties	113
Resilient connection resource usage	115
Specifying resilient connections	116
Seeding the fingerprint cache for remote client-side deduplication	117
Adding a deduplication load balancing server	117
About the pd.conf configuration file for NetBackup deduplication	119
pd.conf file settings for NetBackup deduplication	119
Editing the pd.conf deduplication file	126
About the contentrouter.cfg file for NetBackup deduplication	127
About saving the deduplication storage server configuration	128
Saving the deduplication storage server configuration	129
Editing a deduplication storage server configuration file	129

	Setting the deduplication storage server configuration	131
	About the deduplication host configuration file	132
	Deleting a deduplication host configuration file	132
	Resetting the deduplication registry	132
	Configuring deduplication log file timestamps on Windows	133
	Setting NetBackup configuration options by using bpsetconfig	134
Chapter 6	Monitoring deduplication activity	137
	Monitoring the deduplication rate	137
	Viewing deduplication job details	138
	Deduplication job details	139
	About deduplication storage capacity and usage reporting	141
	About deduplication container files	143
	Viewing storage usage within deduplication container files	143
	Viewing disk reports	145
	Monitoring deduplication processes	146
	Reporting on Auto Image Replication jobs	147
Chapter 7	Managing deduplication	149
	Managing deduplication servers	149
	Viewing deduplication storage servers	150
	Determining the deduplication storage server state	150
	Viewing deduplication storage server attributes	151
	Setting deduplication storage server attributes	152
	Changing deduplication storage server properties	153
	Clearing deduplication storage server attributes	154
	About changing the deduplication storage server name or storage path	155
	Changing the deduplication storage server name or storage path	155
	Removing a load balancing server	157
	Deleting a deduplication storage server	158
	Deleting the deduplication storage server configuration	159
	About shared memory on Windows deduplication storage servers	159
	Managing NetBackup Deduplication Engine credentials	160
	Determining which media servers have deduplication credentials	160
	Adding NetBackup Deduplication Engine credentials	160
	Changing NetBackup Deduplication Engine credentials	161
	Deleting credentials from a load balancing server	161
	Managing deduplication disk pools	161

Viewing deduplication disk pools	162
Determining the deduplication disk pool state	162
Changing the deduplication disk pool state	162
Viewing deduplication disk pool attributes	163
Setting deduplication disk pool attributes	164
Changing deduplication disk pool properties	165
Clearing deduplication disk pool attributes	170
Determining the deduplication disk volume state	171
Changing the deduplication disk volume state	171
Deleting a deduplication disk pool	172
Deleting backup images	173
Disabling client-side deduplication for a client	173
About deduplication queue processing	174
Processing the deduplication transaction queue manually	174
About deduplication data integrity checking	175
Configuring deduplication data integrity checking behavior	176
Deduplication data integrity checking configuration	
settings	178
About managing storage read performance	179
About deduplication storage rebasing	180
Resizing the deduplication storage partition	182
About restoring files at a remote site	183
About restoring from a backup at a target master domain	183
Specifying the restore server	184
Chapter 8	
Troubleshooting	187
About deduplication logs	187
About VxUL logs for deduplication	190
Troubleshooting installation issues	191
Installation on SUSE Linux fails	191
Troubleshooting configuration issues	191
Storage server configuration fails	192
Database system error (220)	192
Server not found error	193
License information failure during configuration	193
The disk pool wizard does not display a volume	194
Troubleshooting operational issues	194
Verify that the server has sufficient memory	195
Backup or duplication jobs fail	195
Client deduplication fails	196
Volume state changes to DOWN when volume is	
unmounted	197

	Errors, delayed response, hangs	198
	Cannot delete a disk pool	198
	Media open error (83)	199
	Media write error (84)	200
	Storage full conditions	201
	Viewing disk errors and events	202
	Deduplication event codes and messages	202
Chapter 9	Host replacement, recovery, and uninstallation	207
	Replacing the deduplication storage server host computer	207
	Recovering from a deduplication storage server disk failure	209
	Recovering from a deduplication storage server failure	211
	Recovering the storage server after NetBackup catalog recovery	212
	About uninstalling media server deduplication	213
	Removing media server deduplication	213
Chapter 10	Deduplication architecture	215
	Deduplication storage server components	215
	Media server deduplication process	217
	Deduplication client components	220
	Client-side deduplication backup process	220
	About deduplication fingerprinting	223
	Data removal process	224
Appendix A	NetBackup appliance deduplication	227
	About NetBackup appliance deduplication	227
	About Fibre Channel to a NetBackup 5020 appliance	228
	Enabling Fibre Channel to a NetBackup 5020 appliance	229
	Disabling Fibre Channel to a NetBackup 5020 appliance	230
	Displaying NetBackup 5020 appliance Fibre Channel port information	230
Index	233

Introducing NetBackup deduplication

This chapter includes the following topics:

- [About NetBackup deduplication](#)
- [New features and enhancements for NetBackup 7.5](#)

About NetBackup deduplication

Symantec NetBackup provides the deduplication options that let you deduplicate data everywhere, as close to the source of data as you require.

Deduplication everywhere provides significant return on investment, as follows.

- Reduce the amount of data that is stored.
- Reduce backup bandwidth.
Reduced bandwidth can be especially important when you want to limit the amount of data that a client sends over the network. Over the network can be to a backup server or for image duplication between remote locations.
- Reduce backup windows.
- Reduce infrastructure.

About NetBackup deduplication options

Deduplication everywhere lets you choose at which point in the backup process to perform deduplication. NetBackup can manage your deduplication wherever you implement it in the backup stream.

[Table 1-1](#) describes the options for deduplication.

Table 1-1 NetBackup deduplication options

Type	Description
NetBackup Client Deduplication Option	<p>With NetBackup client-side deduplication, clients deduplicate their backup data and then send it directly to the storage destination. A media server does not deduplicate the data.</p> <p>See “About NetBackup Client Deduplication” on page 28.</p>
NetBackup Media Server Deduplication Option	<p>NetBackup clients send their backups to a NetBackup media server, which deduplicates the backup data. A NetBackup media server hosts the NetBackup Deduplication Engine, which writes the data to the storage and manages the deduplicated data.</p> <p>See “About the NetBackup Media Server Deduplication Option” on page 23.</p>
NetBackup appliance deduplication	<p>Symantec provides a hardware and a software solution that includes NetBackup deduplication. The NetBackup 5200 series of appliances run the SUSE Linux operating system and include NetBackup software and disk storage.</p> <p>The NetBackup appliances have their own documentation set.</p> <p>See “About NetBackup appliance deduplication” on page 227.</p>
PureDisk deduplication	<p>NetBackup PureDisk is a deduplication solution that provides bandwidth-optimized backups of data in remote offices.</p> <p>You use the PureDisk interfaces to install, configure, and manage the PureDisk servers, storage pools, and client backups. You do not use NetBackup to configure or manage the storage or backups.</p> <p>PureDisk has its own documentation set.</p> <p>See the <i>NetBackup PureDisk Getting Started Guide</i>.</p> <p>A PureDisk storage pool can be a storage destination for both the NetBackup Client Deduplication Option and the NetBackup Media Server Deduplication Option.</p>
PureDisk appliance deduplication	<p>Symantec provides a hardware and a software solution that includes PureDisk deduplication. The NetBackup 5000 series of appliances run the PDOS operating system and include PureDisk software and disk storage.</p> <p>The NetBackup appliances have their own documentation set.</p> <p>See “About NetBackup appliance deduplication” on page 227.</p>

Table 1-1 NetBackup deduplication options (*continued*)

Type	Description
Third-party vendor appliance deduplication	<p>The NetBackup OpenStorage option lets third-party vendor appliances function as disk storage for NetBackup.</p> <p>The disk appliance provides the storage and it manages the storage. A disk appliance may provide deduplication functionality. NetBackup backs up and restores client data and manages the life cycles of the data.</p>


How deduplication works

Deduplication is a method of retaining only one unique instance of backup data on storage media. Redundant data is replaced with a pointer to the unique data copy. Deduplication occurs on both a file level and a file segment level. When two or more files are identical, deduplication stores only one copy of the file. When two or more files share identical content, deduplication breaks the files into segments and stores only one copy of each unique file segment.

Deduplication significantly reduces the amount of storage space that is required for the NetBackup backup images.

Figure 1-1 is a diagram of file segments that are deduplicated.

Figure 1-1 File deduplication



The following list describes how NetBackup derives unique segments to store:

- The deduplication engine breaks file 1 into segments A, B, C, D, and E.
- The deduplication engine breaks file 2 into segments A, B, Q, D, and L.
- The deduplication engine stores file segments A, B, C, D, and E from file 1 and file segments Q, and L from file 2. The deduplication engine does not store file

segments A, B, and D from file 2. Instead, it points to the unique data copies of file segments A, B, and D that were already written from file 1.

More detailed information is available.

See [“Media server deduplication process”](#) on page 217.

New features and enhancements for NetBackup 7.5

The following deduplication features and improvements are included in the NetBackup 7.5 release:

- Support for the AIX 5.3, 6.1, and 7.1 operating systems for deduplication servers and for client-side deduplication.
- 64-TB support for media server deduplication pools.
Information about how to upgrade to more than 32 TB of storage space is available.
See [“About support for more than 32-TB of storage”](#) on page 68.
- Resilient network connections provide improved support for remote office client deduplication.
See [“About remote office client deduplication”](#) on page 31.
- iSCSI support.
Information about the requirements for iSCSI support in NetBackup is available.
See [“About deduplication server requirements”](#) on page 26.
PureDisk 6.6.3 supports iSCSI disks in PureDisk storage pools only if the storage pool is deployed exclusively for PureDisk Deduplication Option (PDDO) use. iSCSI storage pools for PDDO use must be configured with the XFS file system and cannot be clustered.
The PureDisk 6.6.3 documentation does not describe how to configure or manage a PureDisk storage pool that includes iSCSI disks. Information about how to use iSCSI disks in a PureDisk environment is in the PureDisk 6.6.1 documentation.
<http://www.symantec.com/docs/DOC3878>
For known issues about iSCSI storage pools, a Symantec tech note is available.
<http://www.symantec.com/docs/TECH137146>
- Enhancements that improve restore and duplication performance.
See [“About managing storage read performance”](#) on page 179.
- Deduplication integrity enhancements.
See [“About deduplication data integrity checking”](#) on page 175.
- Windows storage server performance enhancements.

Interprocess communication changes on Windows hosts improve performance to be similar to UNIX and Linux hosts.

This change affects upgrades to NetBackup 7.5.

See [“About shared memory on Windows deduplication storage servers”](#) on page 159.

- FlashBackup performance improvements.
- Backup image delete and import performance improvements.
- NetBackup now reserves 4 percent of the storage space for the deduplication database and transaction logs rather than 10 percent.
- A new stream handler for EMC NDMP.
- Fibre Channel connections to NetBackup 5020 appliances. Fibre Channel is supported on x86-64 hosts that run the Red Hat Enterprise Linux 5 or SUSE Enterprise Linux Server 10 SP1 operating systems.
See [“About Fibre Channel to a NetBackup 5020 appliance”](#) on page 228.
- The performance of the first backup of a remote client can be improved.
See [“Seeding the fingerprint cache for remote client-side deduplication”](#) on page 117.
- Garbage collection now occurs during regularly scheduled queue processing. Therefore, the `sched_GarbageCollection.log` file no longer exists.

Planning your deployment

This chapter includes the following topics:

- [Planning your deduplication deployment](#)
- [About the deduplication tech note](#)
- [NetBackup naming conventions](#)
- [About the deduplication storage destination](#)
- [About the NetBackup Media Server Deduplication Option](#)
- [About NetBackup Client Deduplication](#)
- [About remote office client deduplication](#)
- [About NetBackup Deduplication Engine credentials](#)
- [About the network interface for deduplication](#)
- [About deduplication port usage](#)
- [About deduplication compression](#)
- [About deduplication encryption](#)
- [About optimized synthetic backups and deduplication](#)
- [About deduplication and SAN Client](#)
- [About optimized duplication and replication](#)
- [About deduplication performance](#)
- [About deduplication stream handlers](#)
- [Deployment best practices](#)

- [Replacing the PureDisk Deduplication Option with Media Server Deduplication on the same host](#)
- [Migrating from PureDisk to the NetBackup Media Server Deduplication option](#)
- [Migrating from another storage type to deduplication](#)

Planning your deduplication deployment

[Table 2-1](#) provides an overview of planning your deployment of NetBackup deduplication.

Table 2-1 Deployment overview

Step	Deployment task	Where to find the information
Step 1	Read the deduplication tech note	See “About the deduplication tech note” on page 21.
Step 2	Determine the storage destination	See “About the deduplication storage destination” on page 22.
Step 3	Determine which type of deduplication to use	See “About the NetBackup Media Server Deduplication Option” on page 23. See “About NetBackup Client Deduplication” on page 28. See “About remote office client deduplication” on page 31.
Step 4	Determine the requirements for deduplication hosts	See “About NetBackup deduplication servers” on page 25. See “About deduplication server requirements” on page 26. See “About client deduplication requirements and limitations” on page 30. See “About the network interface for deduplication” on page 33. See “About deduplication port usage” on page 33. See “About scaling deduplication” on page 55. See “About deduplication performance” on page 52.
Step 5	Determine the credentials for deduplication	See “About NetBackup Deduplication Engine credentials” on page 32.
Step 6	Read about compression and encryption	See “About deduplication compression” on page 34. See “About deduplication encryption” on page 34.
Step 7	Read about optimized synthetic backups	See “About optimized synthetic backups and deduplication” on page 35.

Table 2-1 Deployment overview (*continued*)

Step	Deployment task	Where to find the information
Step 8	Read about deduplication and SAN Client	See “ About deduplication and SAN Client ” on page 36.
Step 9	Read about optimized duplication and replication	See “ About optimized duplication and replication ” on page 36.
Step 10	Read about stream handlers	See “ About deduplication stream handlers ” on page 54.
Step 11	Read about best practices for implementation	See “ Deployment best practices ” on page 54.
Step 12	Determine the storage requirements and provision the storage	See “ About provisioning the deduplication storage ” on page 65. See “ About deduplication storage requirements ” on page 66. See “ About deduplication storage capacity ” on page 68. See “ About the deduplication storage paths ” on page 69.
Step 13	Replace a PDDO host or migrate from PDDO to NetBackup deduplication	See “ Replacing the PureDisk Deduplication Option with Media Server Deduplication on the same host ” on page 60. See “ Migrating from PureDisk to the NetBackup Media Server Deduplication option ” on page 61.
Step 14	Migrate from other storage to NetBackup deduplication	See “ Migrating from another storage type to deduplication ” on page 62.

About the deduplication tech note

Symantec provides a tech note that includes the following:

- Currently supported systems
- Media server and client sizing information
- Configuration, operational, and troubleshooting updates
- And more

See the following link:

<http://symantec.com/docs/TECH77575>

NetBackup naming conventions

NetBackup has rules for naming logical constructs. Generally, names are case sensitive. The following set of characters can be used in user-defined names, such as clients, disk pools, backup policies, storage lifecycle policies, and so on:

- Alphabetic (A-Z a-z) (names are case sensitive)
- Numeric (0-9)
- Period (.)
- Plus (+)
- Minus (-)
Do not use a minus as the first character.
- Underscore (_)

Note: No spaces are only allowed.

The naming conventions for the NetBackup Deduplication Engine differ from these NetBackup naming conventions.

See “[About NetBackup Deduplication Engine credentials](#)” on page 32.

About the deduplication storage destination

Several destinations exist for the NetBackup deduplication, as shown in the following table.

Table 2-2 NetBackup deduplication destinations

Destination	Description
Media Server Deduplication Pool	<p>A NetBackup Media Server Deduplication Pool represents the disk storage that is attached to a NetBackup media server. NetBackup deduplicates the data and hosts the storage.</p> <p>If you use this destination, use this guide to plan, implement, configure, and manage deduplication and the storage. When you configure the storage server, select Media Server Deduplication Pool as the storage type.</p> <p>For a Media Server Deduplication Pool storage destination, all hosts that are used for the deduplication must be NetBackup 7.0 or later. Hosts include the master server, the media servers, and the clients that deduplicate their own data. Integrated deduplication means that the components installed with NetBackup perform deduplication.</p>

Table 2-2 NetBackup deduplication destinations (*continued*)

Destination	Description
PureDisk Deduplication Pool	<p>A NetBackup PureDisk Deduplication Pool represents a PureDisk storage pool. NetBackup deduplicates the data, and the PureDisk environment hosts the storage.</p> <p>If you use a PureDisk Deduplication Pool, use the following documentation:</p> <ul style="list-style-type: none"> ■ The PureDisk documentation to plan, implement, configure, and manage the PureDisk environment, which includes the storage. See the <i>NetBackup PureDisk Getting Started Guide</i>. ■ This guide to configure backups and deduplication in NetBackup. When you configure the storage server, select PureDisk Deduplication Pool as the storage type. <p>A PureDisk Deduplication Pool destination requires that PureDisk be at release 6.6 or later.</p> <p>A NetBackup 5000 series appliance also provides a PureDisk storage pool to which NetBackup can send deduplicated data.</p>

In addition to the previous storage destinations, the PureDisk Deduplication Option provides deduplication storage for a NetBackup environment. The PureDisk Storage Pool Authority provides an agent that you install on a NetBackup media server. This solution was created before NetBackup included integrated deduplication. If you use this storage option, use the PureDisk documentation to plan, implement, configure, and manage the storage and to configure NetBackup to use the agent. For a PureDisk storage pool destination, you can use NetBackup 6.5 or later NetBackup hosts. Hosts include the master server and the media servers.

About the NetBackup Media Server Deduplication Option

The NetBackup Media Server Deduplication Option exists in the Symantec OpenStorage framework. A storage server writes data to the storage and reads data from the storage; the storage server must be a NetBackup media server. The storage server hosts the core components of deduplication. The storage server also deduplicates the backup data. It is known as a deduplication storage server.

For a backup, the NetBackup client software creates the image of backed up files as for a normal backup. The client sends the backup image to the deduplication storage server, which deduplicates the data. The deduplication storage server writes the data to disk.

See [“About NetBackup deduplication servers”](#) on page 25.

The NetBackup Media Server Deduplication Option is integrated into NetBackup. It uses the NetBackup administration interfaces, commands, and processes for configuring and executing backups and for configuring and managing the storage. Deduplication occurs when NetBackup backs up a client to a deduplication storage destination. You do not have to use the separate PureDisk interfaces to configure and use deduplication.

The NetBackup Media Server Deduplication Option integrates with NetBackup application agents that are optimized for the client stream format. Agents include but are not limited to Microsoft Exchange and Microsoft SharePoint Agents.

Figure 2-1 shows NetBackup media server deduplication. The deduplication storage server is a media server on which the deduplication core components are enabled. The storage destination is a **Media Server Deduplication Pool**.

Figure 2-1 NetBackup media server deduplication

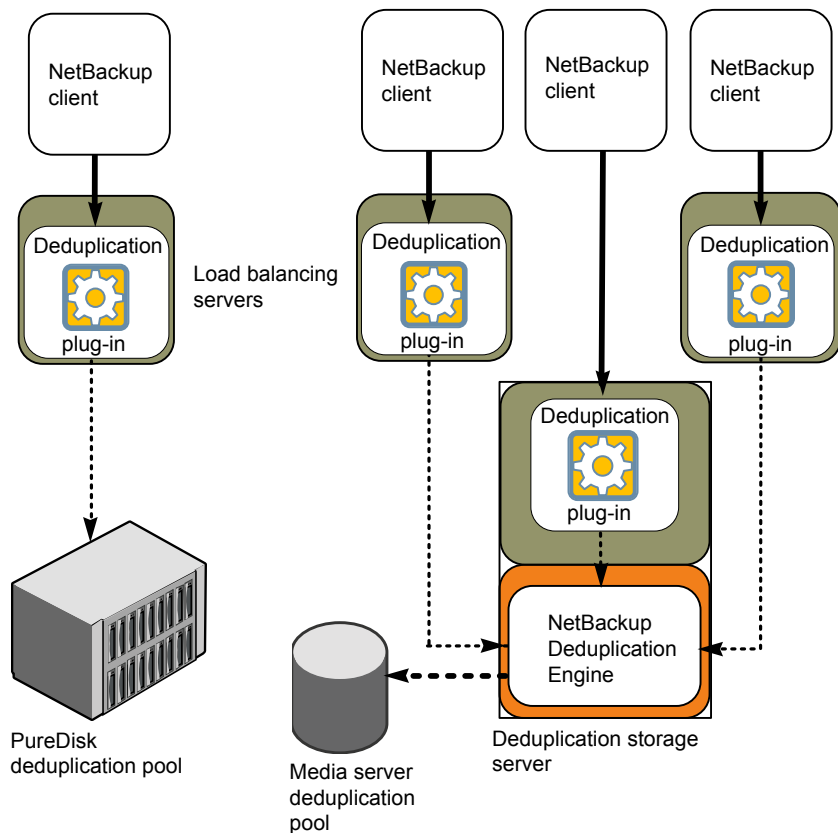


Table 2-3 NetBackup deduplication servers (*continued*)

Host	Description
Load balancing server	<p>You can configure other NetBackup media servers to help deduplicate data. They perform file fingerprint calculations for deduplication, and they send the unique results to the storage server. These helper media servers are called load balancing servers.</p> <p>See “About deduplication fingerprinting” on page 223.</p> <p>A NetBackup media server becomes a load balancing server when two things occur:</p> <ul style="list-style-type: none"> ■ You enable the media server for deduplication load balancing duties. You do so when you configure the storage server or later by modifying the storage server properties. ■ You select it in the storage unit for the deduplication pool. <p>See “Introduce load balancing servers gradually” on page 56.</p> <p>Load balancing servers also perform restore and duplication jobs.</p> <p>Load balancing servers can be any supported server type for deduplication. They do not have to be the same type as the storage server.</p>

About deduplication nodes

A media server deduplication node is a deduplication storage server, load balancing servers (if any), the clients that are backed up, and the storage. Each node manages its own storage. Deduplication within each node is supported; deduplication between nodes is not supported.

Multiple media server deduplication nodes can exist. Nodes cannot share servers, storage, or clients.

About deduplication server requirements

The host computer’s CPU and memory constrain how many jobs can run concurrently. The storage server requires enough capability for deduplication and for storage management unless you offload some of the deduplication to load balancing servers.

[Table 2-4](#) shows the minimum requirements for deduplication servers. NetBackup deduplication servers are always NetBackup media servers.

Processors for deduplication should have a high clock rate and high floating point performance. Furthermore, high throughput per core is desirable. Each backup stream uses a separate core.

Intel and AMD have similar performance and perform well on single core throughput.

Newer SPARC processors, such as the SPARC64 VII, provide the single core throughput that is similar to AMD and Intel. Alternatively, UltraSPARC T1 and T2 single core performance does not approach that of the AMD and Intel processors. Tests show that the UltraSPARC processors can achieve high aggregate throughput. However, they require eight times as many backup streams as AMD and Intel processors to do so.

Table 2-4 Deduplication server minimum requirements

Component	Storage server	Load balancing server or PureDisk Deduplication Option host
CPU	<p>Symantec recommends at least a 2.2-GHz clock rate. A 64-bit processor is required.</p> <p>At least four cores are required. Symantec recommends eight cores.</p> <p>For 64 TBs of storage, Intel x86-64 architecture requires eight cores.</p>	<p>Symantec recommends at least a 2.2-GHz clock rate. A 64-bit processor is required.</p> <p>At least two cores are required. Depending on throughput requirements, more cores may be helpful.</p>
RAM	<p>4 GBs to 64 GBs.</p> <p>If your storage exceeds 4 TBs, Symantec recommends at least 1 GB more of memory for every terabyte of additional storage. For example, 10 TBs of back-end data require 10 GBs of RAM, 32 TBs require 32 GBs of RAM, and so on.</p> <p>For 64 TBs of storage, 64 GBs of RAM are required.</p>	<p>4 GBs.</p>
Operating system	<p>The operating system must be a supported 64-bit operating system.</p> <p>See the operating system compatibility list for your NetBackup release on the NetBackup Enterprise Server landing page on the Symantec Support Web site.</p>	<p>The operating system must be a supported 64-bit operating system.</p> <p>See the operating system compatibility list for your NetBackup release on the NetBackup Enterprise Server landing page on the Symantec Support Web site.</p>

A deduplication tech note provides detailed information about and examples for sizing the hosts for deduplication. Information includes the number of NICs or HBAs per server that are required to support your performance objectives.

See “[About the deduplication tech note](#)” on page 21.

Note: In some environments, a single host can function as both a NetBackup master server and as a deduplication server. Such environments typically run fewer than 100 total backup jobs a day. (Total backup jobs means backups to any storage destination, including deduplication and nondeduplication storage.) If you perform more than 100 backups a day, deduplication operations may affect master server operations.

See [“About deduplication performance”](#) on page 52.

See [“About deduplication queue processing”](#) on page 174.

About media server deduplication limitations

NetBackup media server deduplication and Symantec Backup Exec deduplication cannot reside on the same host. If you use both NetBackup and Backup Exec deduplication, each product must reside on a separate host.

NetBackup deduplication components cannot reside on the same host as the PureDisk Deduplication Option (PDDO) agent that is installed from the PureDisk distribution.

You cannot upgrade to NetBackup 7.0 or later a NetBackup media server that hosts a PDDO agent. If the NetBackup 7.0 installation detects the PDDO agent, the installation fails. To upgrade a NetBackup media server that hosts a PDDO agent, you must first remove the PDDO agent. You then can use that host as a front end for your PureDisk Storage Pool Authority. (The host must be a host type that is supported for NetBackup deduplication.)

See [“Replacing the PureDisk Deduplication Option with Media Server Deduplication on the same host”](#) on page 60.

See the *NetBackup PureDisk Deduplication Option (PDDO) Guide*.

Deduplication within each media server deduplication node is supported; global deduplication between nodes is not supported.

About NetBackup Client Deduplication

With normal deduplication, the client sends the full backup data stream to the media server. The deduplication engine on the media server processes the stream, saving only the unique segments.

With NetBackup Client Deduplication, the client hosts the deduplication plug-in that duplicates the backup data. The NetBackup client software creates the image of backed up files as for a normal backup. Next, the deduplication plug-in breaks the backup image into segments and compares them to all of the segments that

are stored in that deduplication node. The plug-in then sends only the unique segments to the NetBackup Deduplication Engine on the storage server. The engine writes the data to a media server deduplication pool.

Client deduplication does the following:

- Reduces network traffic. The client sends only unique file segments to the storage server. Duplicate data is not sent over the network.
- Distributes some deduplication processing load from the storage server to clients. (NetBackup does not balance load between clients; each client deduplicates its own data.)


NetBackup Client Deduplication is a solution for the following cases:

- Remote office or branch office backups to the data center.
 NetBackup provides resilient network connections for remote office backups. See [“About remote office client deduplication”](#) on page 31.
- LAN connected file server
- Virtual machine backups.

Client-side deduplication is also a useful solution if a client host has unused CPU cycles or if the storage server or load balancing servers are overloaded.

[Figure 2-2](#) shows client deduplication. The deduplication storage server is a media server on which the deduplication core components are enabled. The storage destination is a **Media Server Deduplication Pool**

Figure 2-2 NetBackup client deduplication



A PureDisk storage pool may also be the storage destination. See [“About the deduplication storage destination”](#) on page 22. More information is available. See [“About remote office client deduplication”](#) on page 31. See [“Deduplication client components”](#) on page 220. See [“Client-side deduplication backup process”](#) on page 220.

About client deduplication requirements and limitations

The clients that deduplicate their own data must be at the same revision level as the NetBackup deduplication servers.

For supported systems, see the *NetBackup Release Notes*.

Client deduplication does not support multiple copies per job configured in a NetBackup backup policy. For the jobs that specify multiple copies, the backup images are sent to the storage server and may be deduplicated there.

About remote office client deduplication

WAN backups require more time than local backups in your own domain. WAN backups have an increased risk of failure when compared to local backups. To help facilitate WAN backups, NetBackup provides the capability for resilient network connections. A resilient connection allows backup and restore traffic between a client and NetBackup media servers to function effectively in high-latency, low-bandwidth networks such as WANs.

The use case that benefits the most from resilient connections is client-side deduplication at a remote office that does not have local backup storage. The following items describe the advantages:

- Client deduplication reduces the time that is required for WAN backups by reducing the amount of data that must be transferred.
- The resilient connections provide automatic recovery from network failures and latency (within the parameters from which NetBackup can recover).

When you configure a resilient connection, NetBackup uses that connection for the backups. Use the NetBackup **Resilient Network** host properties to configure NetBackup to use resilient network connections.

See [“Resilient Network properties”](#) on page 113.

See [“Specifying resilient connections”](#) on page 116.

You can improve the performance of the first backup for a remote client.

See [“Seeding the fingerprint cache for remote client-side deduplication”](#) on page 117.

About remote client deduplication data security

Resilient connection traffic is not encrypted. The NetBackup deduplication process can encrypt the data before it is transmitted over the WAN. Symantec recommends that you use the deduplication encryption to protect your data during your remote client backups.

See [“About deduplication encryption”](#) on page 34.

NetBackup does not encrypt the data during a restore job. Therefore, Symantec recommends that you restore data to the original remote client over a private network.

See [“How deduplication restores work”](#) on page 59.

About remote client backup scheduling

NetBackup backup policies use the time zone of the master server for scheduling jobs. If your remote clients are in a different a time zone than your NetBackup master server, you must compensate for the difference. For example, suppose the master server is in Finland (UTC+2) and the remote client is in London (UTC+0). If the backup policy has a window from 6pm to 6am, backups can begin at 4pm on the client. To compensate, you should set the backup window from 8pm to 8am. Alternatively, it may be advisable to use a separate backup policy for each time zone in which remote clients reside.

About NetBackup Deduplication Engine credentials

The NetBackup Deduplication Engine requires credentials. The deduplication components use the credentials when they communicate with the NetBackup Deduplication Engine. The credentials are for the deduplication engine, not for the host on which it runs.

You enter the NetBackup Deduplication Engine credentials when you configure the storage server.

The following are the rules for the credentials:

- For user names and passwords, you can use characters in the printable ASCII range (0x20-0x7E) except for the following characters:
 - Asterisk (*)
 - Backward slash (\) and forward slash (/)
 - Double quote (")
 - Left parenthesis [(] and right parenthesis [)]
- The user name and the password can be up to 63 characters in length.
- Leading and trailing spaces and quotes are ignored.
- The user name and password cannot be empty or all spaces.

Note: Record and save the credentials in case you need them in the future.

Caution: You cannot change the NetBackup Deduplication Engine credentials after you enter them. Therefore, carefully choose and enter your credentials. If you must change the credentials, contact your Symantec support representative.

About the network interface for deduplication

If the deduplication storage server host has more than one network interface, by default the host operating system determines which network interface to use. However, you can specify which interface NetBackup should use for the backup and restore traffic.

To use a specific interface, enter that interface name when you configure the deduplication storage server.

Caution: Carefully enter the network interface. If you make a mistake, the process to recover is time consuming.

See [“Changing the deduplication storage server name or storage path”](#) on page 155.

The NetBackup `REQUIRED_INTERFACE` setting does not affect deduplication processes.

About deduplication port usage

The following table shows the ports that are used for NetBackup deduplication. If firewalls exist between the various deduplication hosts, open the indicated ports on the deduplication hosts. Deduplication hosts are the deduplication storage server, the load balancing servers, and the clients that deduplicate their own data.

If you have only a storage server and no load balancing servers or clients that deduplicate their own data, you do not have to open firewall ports.

Table 2-5 Deduplication ports

Port	Usage
10082	The NetBackup Deduplication Engine (<code>spoold</code>). Open this port between the hosts that deduplicate data.
10085	The deduplication database (<code>postgres</code>). The connection is internal to the storage server, from <code>spad</code> to <code>spoold</code> . You do not have to open this port.
10102	The NetBackup Deduplication Manager (<code>spad</code>). Open this port between the hosts that deduplicate data.

About deduplication compression

NetBackup provides compression for the deduplicated data. It is separate from and different than NetBackup policy-based compression. By default, deduplication compression is enabled.

Symantec recommends that you use the deduplication compression.

Note: Do not enable compression by selecting the **Compression** setting on the **Attributes** tab of the **Policy** dialog box. If you do, NetBackup compresses the data before it reaches the deduplication plug-in that deduplicates it. Consequently, deduplication rates are very low.

See [“Use deduplication compression and encryption”](#) on page 57.

See [“About the pd.conf configuration file for NetBackup deduplication”](#) on page 119.

See [“pd.conf file settings for NetBackup deduplication ”](#) on page 119.

About deduplication encryption

NetBackup provides encryption for the deduplicated data. It is separate from and different than NetBackup policy-based encryption. By default, deduplication encryption is disabled.

Symantec recommends that you use deduplication encryption.

The following is the behavior for the encryption that occurs during the deduplication process:

- If you enable encryption on a client that deduplicates its own data, the client encrypts the data before it sends it to the storage server. The data remains encrypted on the storage.
Data also is transferred from the client over a Secure Sockets Layer to the server regardless of whether or not the data is encrypted. Therefore, data transfer from the clients that do not deduplicate their own data is also protected.
- If you enable encryption on a load balancing server, the load balancing server encrypts the data. It remains encrypted on storage.
- If you enable encryption on the storage server, the storage server encrypts the data. It remains encrypted on storage. If the data is already encrypted, the storage server does not encrypt it.

Deduplication uses the Blowfish algorithm for encryption.

Note: Do not enable encryption by selecting the **Encryption** setting on the **Attributes** tab of the **Policy** dialog box. If you do, NetBackup encrypts the data before it reaches the deduplication plug-in that deduplicates it. Consequently, deduplication rates are very low.

See [“Use deduplication compression and encryption”](#) on page 57.

See [“Enabling deduplication encryption”](#) on page 87.

See [“About the pd.conf configuration file for NetBackup deduplication”](#) on page 119.

See [“pd.conf file settings for NetBackup deduplication ”](#) on page 119.

About optimized synthetic backups and deduplication

Optimized synthetic backups are a more efficient form of synthetic backup. A media server uses messages to instruct the storage server which full and incremental backup images to use to create the synthetic backup. The storage server constructs (or synthesizes) the backup image directly on the disk storage. Optimized synthetic backups require no data movement across the network..

Optimized synthetic backups are faster than a synthetic backup. Regular synthetic backups are constructed on the media server. They are moved across the network from the storage server to the media server and synthesized into one image. The synthetic image is then moved back to the storage server.

The target storage unit's deduplication pool must be the same deduplication pool on which the source images reside.

See [“Configuring optimized synthetic backups for deduplication”](#) on page 89.

If NetBackup cannot produce the optimized synthetic backup, NetBackup creates the more data-movement intensive synthetic backup.

In NetBackup, the **OptimizedImage** attribute enables optimized synthetic backups. It applies to both storage servers and deduplication pools. Beginning with NetBackup 7.1, the **OptimizedImage** attribute is enabled by default on storage servers and media server deduplication pools. For the storage servers and the disk pools that you created in NetBackup releases earlier than 7.1, you must set the **OptimizedImage** attribute on them so they support optimized synthetic backups.

See [“Setting deduplication storage server attributes”](#) on page 152.

See [“Setting deduplication disk pool attributes”](#) on page 164.

About deduplication and SAN Client

SAN Client is a NetBackup optional feature that provides high speed backups and restores of NetBackup clients. Fibre Transport is the name of the NetBackup high-speed data transport method that is part of the SAN Client feature. The backup and restore traffic occurs over a SAN.

SAN clients can be used with the deduplication option; however, the deduplication must occur on the media server, not the client. Configure the media server to be both a deduplication storage server (or load balancing server) and an FT media server. The SAN client backups are then sent over the SAN to the deduplication server/FT media server host. At that media server, the backup stream is deduplicated.

Do not enable client deduplication on SAN Clients. The data processing for deduplication is incompatible with the high-speed transport method of Fibre Transport. Client-side deduplication relies on two-way communication over the LAN with the media server. A SAN client streams the data to the FT media server at a high rate over the SAN.

About optimized duplication and replication

NetBackup supports several methods for optimized duplication and replication of deduplicated data.

The following table lists the duplication methods NetBackup supports between media server deduplication pools.

Table 2-6 NetBackup OpenStorage optimized duplication and replication methods

Optimized duplication method	Description
Within the same NetBackup domain	See “About MSDP optimized duplication within the same domain” on page 36.
To a remote NetBackup domain	See “About NetBackup Auto Image Replication” on page 47.

About MSDP optimized duplication within the same domain

Optimized duplication within the same domain copies the backup images from one **Media Server Deduplication Pool** to a **Media Server Deduplication Pool** in the same domain. The source and the destination must use the same NetBackup master server. The optimized duplication operation is more efficient than normal

duplication. Only the unique, deduplicated data segments are transferred. Optimized duplication reduces the amount of data that is transmitted over your network.

Optimized duplication is a good method to copy your backup images off-site for disaster recovery.

The following sections provide the conceptual information about optimized duplication.

A process topic describes the configuration process.

See [“Configuring optimized duplication of deduplicated data”](#) on page 94.

Optimized MSDP duplication within the same domain requirements

The following are the requirements for optimized duplication within the same NetBackup domain:

- If the source images reside on a NetBackup **Media Server Deduplication Pool**, the destination can be another **Media Server Deduplication Pool** or a **PureDisk Deduplication Pool**. (In NetBackup, a PureDisk storage pool is configured as a **PureDisk Deduplication Pool**.)
 If the destination is a PureDisk storage pool, the PureDisk environment must be at release level 6.6 or later.
- If the source images reside on a PureDisk storage pool, the destination must be another PureDisk storage pool.
 Both PureDisk environments must be at release level 6.6 or later.
- The source storage and the destination storage must have at least one media server in common.
 See [“About the media servers for optimized MSDP duplication within the same domain”](#) on page 38.
- In the storage unit you use for the destination for the optimized duplication, you must select only the common media server or media servers.
 If you select more than one, NetBackup assigns the duplication job to the least busy media server. If you select a media server or servers that are not common, the optimized duplication job fails.
 For more information about media server load balancing, see the *NetBackup Administrator's Guide for UNIX and Linux, Volume I* or the *NetBackup Administrator's Guide for Windows, Volume I*.
- The destination storage unit cannot be the same as the source storage unit.

Optimized MSDP duplication within the same domain limitations

The following are limitations for optimized duplication within the same NetBackup domain:

- You cannot use optimized duplication from a PureDisk storage pool (a **PureDisk Deduplication Pool**) to a **Media Server Deduplication Pool**.
- If an optimized duplication job fails after the configured number of retries, NetBackup does not run the job again.
By default, NetBackup retries an optimized duplication job three times. You can change the number of retries.
See [“Configuring MSDP optimized duplication copy behavior”](#) on page 91.
- Optimized duplication does not work with storage unit groups. If you use a storage unit group as a destination for optimized duplication, NetBackup uses regular duplication.
- Optimized duplication does not support multiple copies. If NetBackup is configured to make multiple new copies from the (source) copy of the backup image, the following occurs:
 - In a storage lifecycle policy, one duplication job creates one optimized duplication copy. If multiple optimized duplication destinations exist, a separate job exists for each destination. This behavior assumes that the device for the optimized duplication destination is compatible with the device on which the source image resides.
If multiple remaining copies are configured to go to devices that are not optimized duplication capable, NetBackup uses normal duplication. One duplication job creates those multiple copies.
 - For other duplication methods, NetBackup uses normal duplication. One duplication job creates all of the copies simultaneously. The other duplication methods include the following: NetBackup Vault, the `bpduplicate` command line, and the duplication option of the **Catalog** utility in the **NetBackup Administration Console**.

See [“Optimized MSDP duplication within the same domain requirements”](#) on page 37.

About the media servers for optimized MSDP duplication within the same domain

For optimized **Media Server Deduplication Pool** duplication within the same domain, the source storage and the destination storage must have at least one media server in common. The common server initiates, monitors, and verifies the

copy operation. The common server requires credentials for both the source storage and the destination storage. (For deduplication, the credentials are for the NetBackup Deduplication Engine, not for the host on which it runs.)

Which server initiates the duplication operation determines if it is a push or a pull operation. If it is physically in the source domain, it is push duplication. If it is in the destination domain, it is a pull duplication. Technically, no advantage exists with a push duplication or a pull duplication. However, the media server that initiates the duplication operation also becomes the write host for the new image copies.

A storage server or a load balancing server can be the common server. The common server must have the credentials and the connectivity for both the source storage and the destination storage.

About MSDP push duplication within the same domain

Figure 2-3 shows a push configuration for optimized within the same domain. The local deduplication node contains normal backups; the remote deduplication node is the destination for the optimized duplication copies. Load balancing server LB_L2 has credentials for both storage servers; it is the common server.

Figure 2-3 Push duplication environment

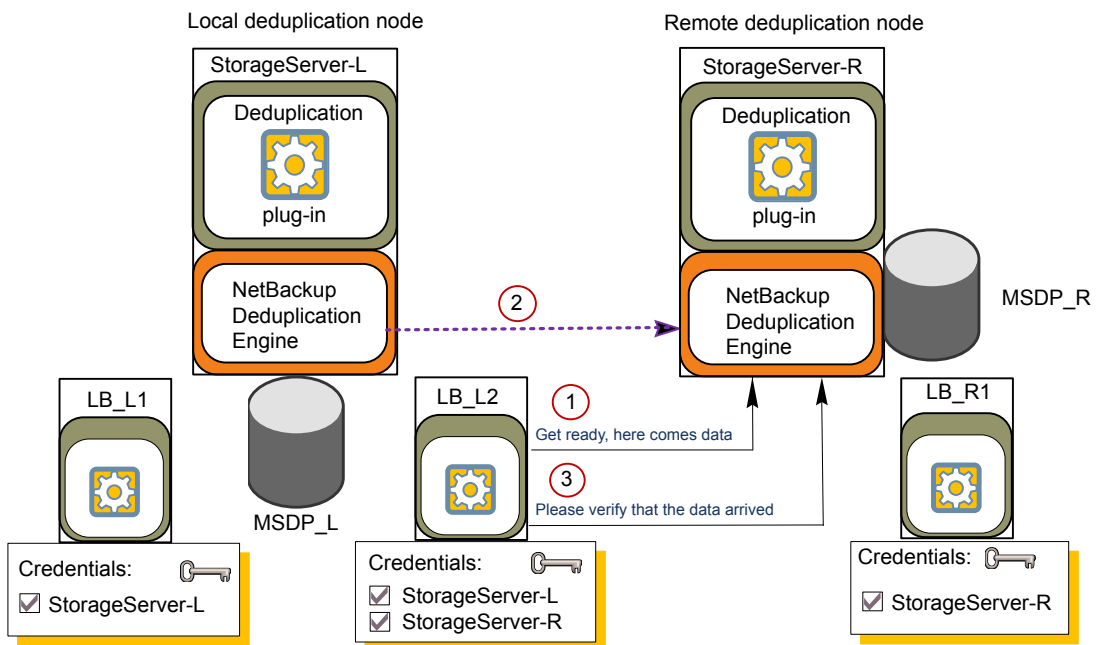


Figure 2-4 shows the settings for the storage unit for the normal backups for the local deduplication node. The disk pool is the **MSDP_L** in the local environment. Because all hosts in the local node are co-located, you can use any available media server for the backups.

Figure 2-4 Storage unit settings for backups to MSDP_L

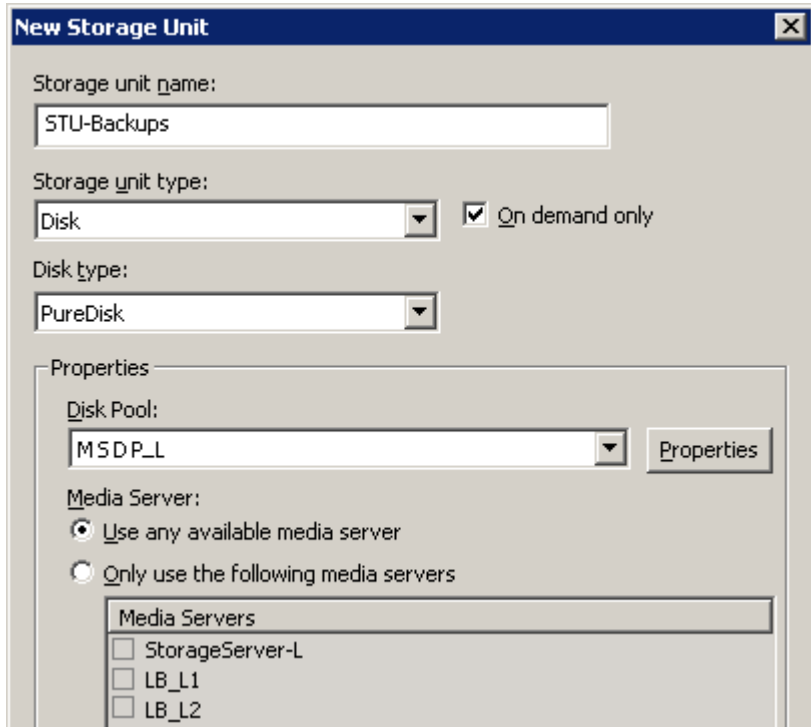
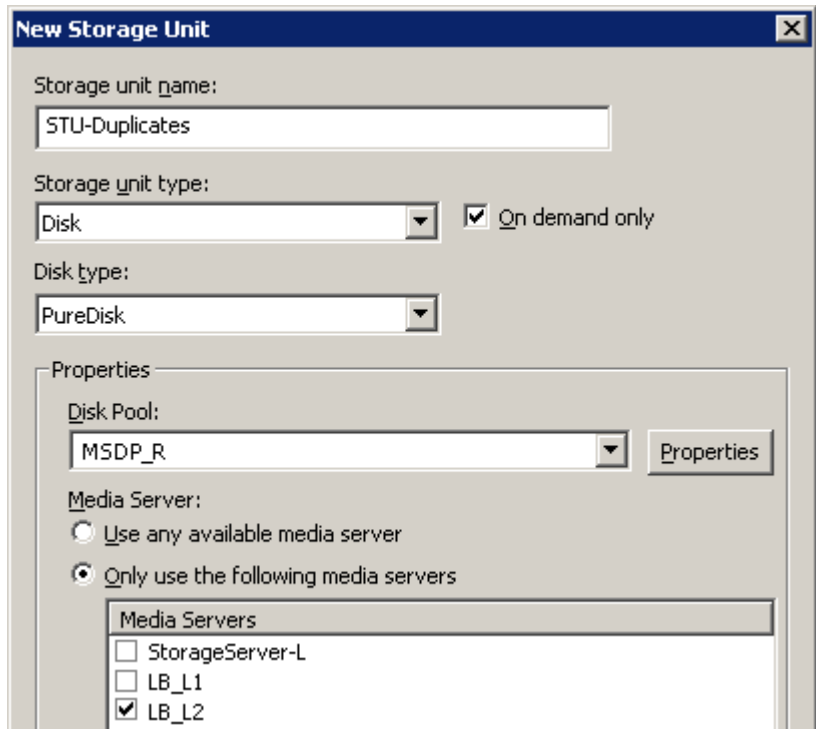


Figure 2-5 shows the storage unit settings for the optimized duplication. The destination is the **MSDP_R** in the remote environment. You must select the common server, so only load balancing server **LB_L2** is selected.

Figure 2-5 Storage unit settings for duplication to MSDP_R



If you use the remote node for backups also, select StorageServer-R and load balancing server LB_R1 in the storage unit for the remote node backups. If you select server LB_L2, it becomes a load balancing server for the remote deduplication pool. In such a case, data travels across your WAN.

Figure 2-6 shows a push duplication from a **Media Server Deduplication Pool** to a PureDisk storage pool. The **Media Server Deduplication Pool** contains normal backups; the **PureDisk Deduplication Pool** is the destination for the optimized duplication copies. The StorageServer-A has credentials for both environments; it is the common media server.

Figure 2-6 Push duplication to a PureDisk storage pool

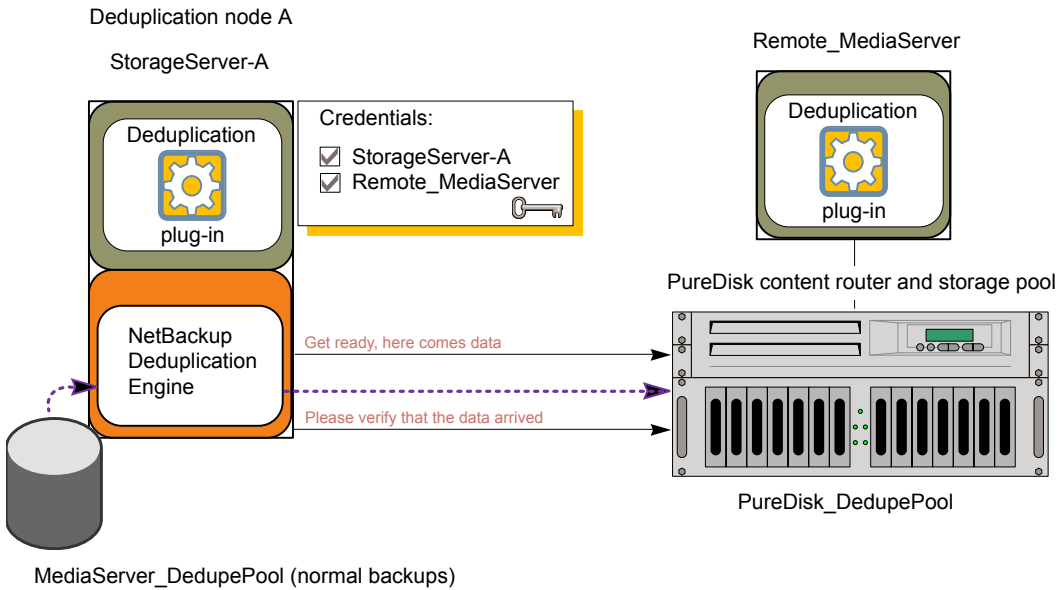


Figure 2-7 shows the storage unit settings for normal backups for the environment in Figure 2-6. The disk pool is the **MediaServer_DedupePool** in the local environment. For normal backups, you do not want a remote host deduplicating data, so only the local host is selected.

Figure 2-7 Storage unit settings for backups to MediaServer_DedupePool

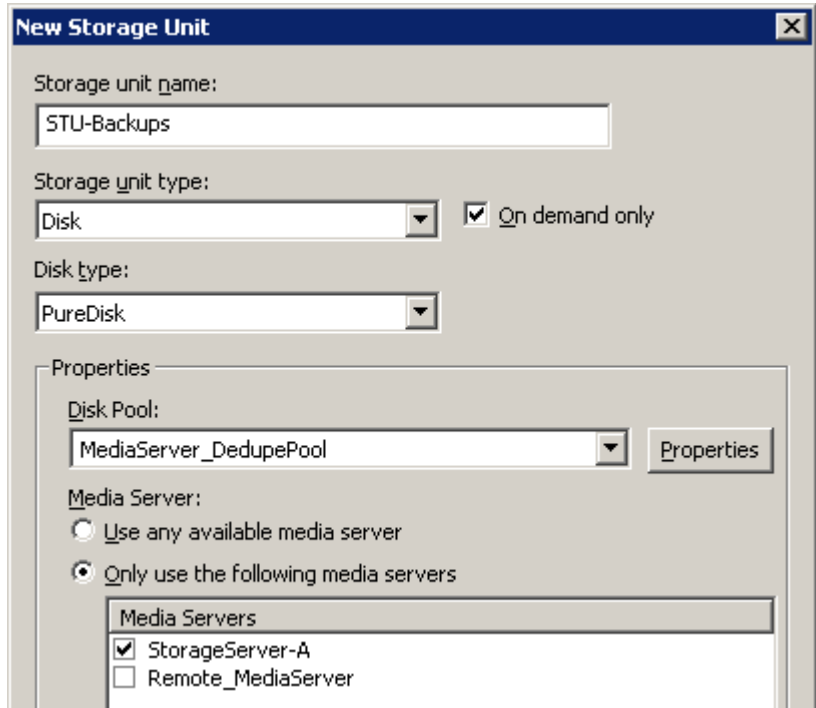


Figure 2-8 shows the storage unit settings for duplication for the environment in Figure 2-6. The disk pool is the **PureDisk_DedupePool** in the remote environment. You must select the common server, so only the local media server is selected. If this configuration were a pull configuration, the remote host would be selected in the storage unit.

Figure 2-8 Storage unit settings for duplication to PureDisk_DedupePool

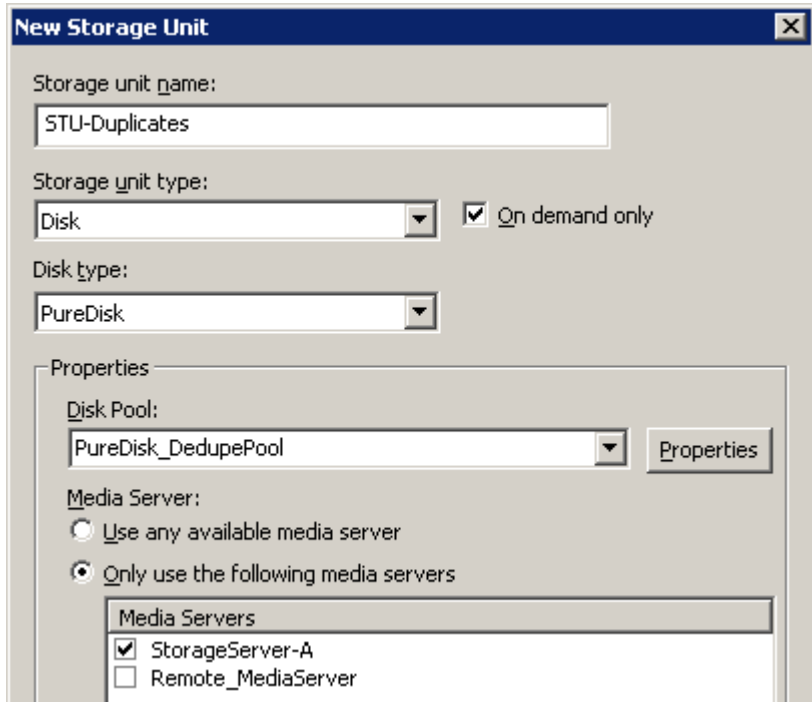
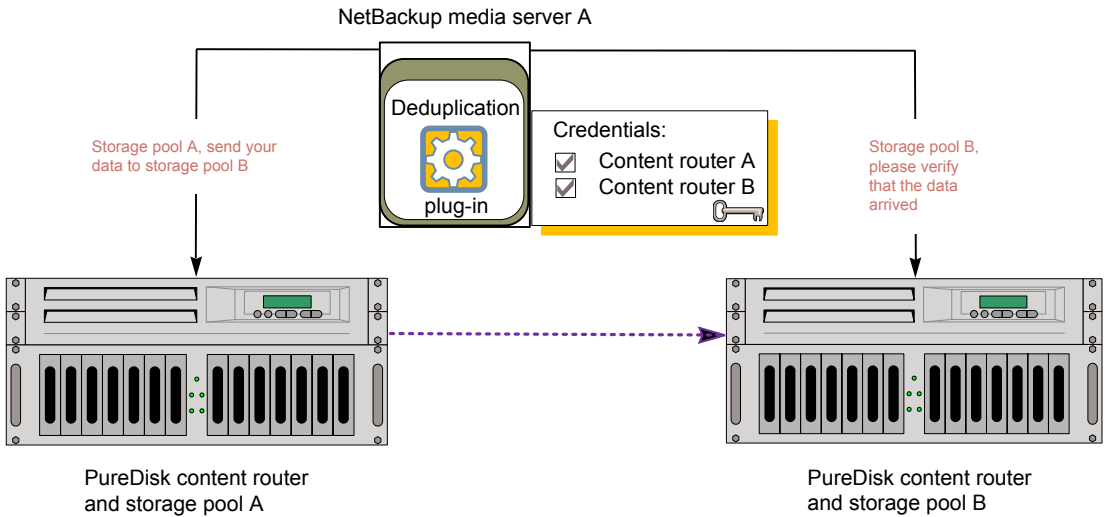


Figure 2-9 shows optimized duplication between two PureDisk storage pools. NetBackup media server A has credentials for both storage pools; it initiates, monitors, and verifies the optimized duplication. In the destination storage unit, the common server (media server A) is selected. This configuration is a push configuration. For a **PureDisk Deduplication Pool** (that is, a PureDisk storage pool), the PureDisk content router functions as the storage server.

Figure 2-9 Storage pool duplication



You can use a load balancing server when you duplicate between two NetBackup deduplication pools. However, it is more common between two PureDisk storage pools.

About MSDP pull duplication within the same domain

Figure 2-10 shows a pull configuration for optimized duplication within the same domain. Deduplication node A contains normal backups; deduplication node B is the destination for the optimized duplication copies. Host B has credentials for both nodes; it is the common server.

Figure 2-10 Pull duplication

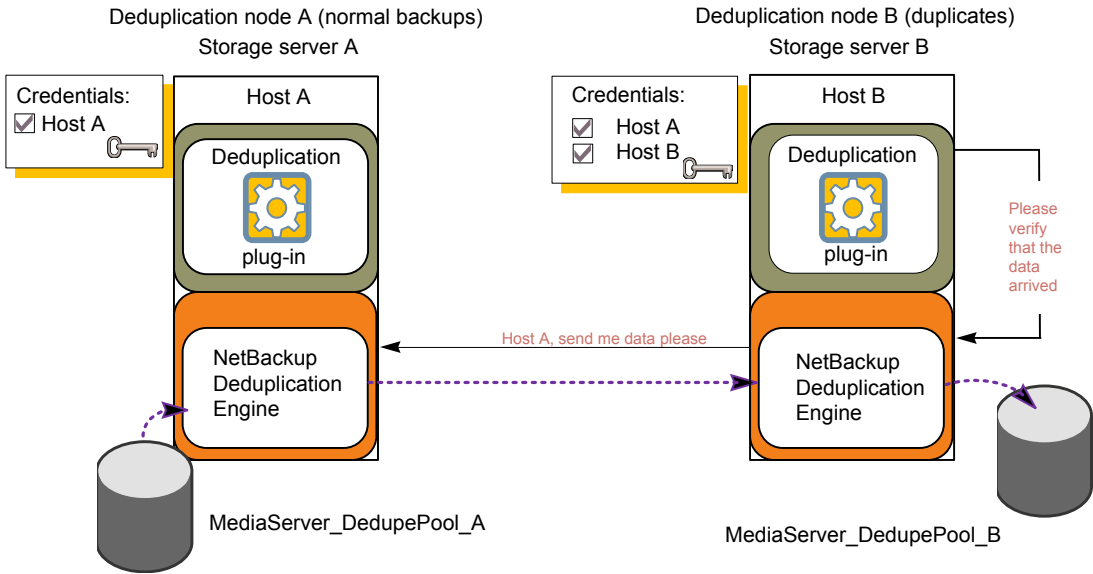
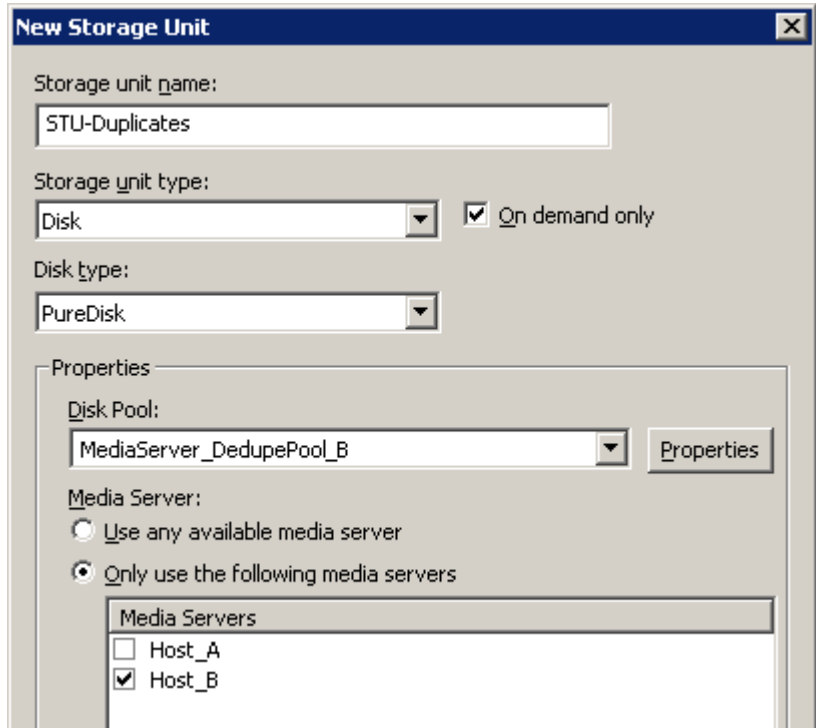


Figure 2-11 shows the storage unit settings for the duplication destination. They are similar to the push example except host B is selected. Host B is the common server, so it must be selected in the storage unit.

Figure 2-11 Pull duplication storage unit settings



If you use node B for backups also, select host B and not host A in the storage unit for the node B backups. If you select host A, it becomes a load balancing server for the node B deduplication pool.

About NetBackup Auto Image Replication

With NetBackup media server deduplication, you can automatically replicate backup images to a different master server domain. This process is referred to as Auto Image Replication.

Auto Image Replication is different than NetBackup optimized duplication within the same domain.

See [“About MSDP optimized duplication within the same domain”](#) on page 36.

The ability to replicate backups to storage in other NetBackup domains, often across various geographical sites, helps facilitate the following disaster recovery needs:

- One-to-one model

A single production datacenter can back up to a disaster recovery site.

- **One-to-many model**

A single production datacenter can back up to multiple disaster recovery sites. See [“One-to-many Auto Image Replication model”](#) on page 49.

- **Many-to-one model**

Remote offices in multiple domains can back up to a storage device in a single domain.

- **Many-to-many model**

Remote datacenters in multiple domains can back up multiple disaster recovery sites.

Note: Although Auto Image Replication is a disaster recovery solution, the administrator cannot directly restore to clients in the primary (or originating) domain from the target master domain.

[Table 2-7](#) is an overview of the process, generally describing the events in the originating and target domains.

Table 2-7 Auto Image Replication process overview

Event	Domain in which event occurs	Event description
1	Originating master (Domain 1)	<p>Clients are backed up according to a policy that indicates a storage lifecycle policy as the Policy storage selection.</p> <p>At least one of the operations in the SLP must be configured for replication to a Media Server Deduplication Pool (MSDP) on a target master.</p> <p>See “About the storage lifecycle policies required for Auto Image Replication” on page 102.</p>
2	Target master (Domain 2)	The deduplication storage server in the target domain recognizes that a replication event has occurred and notifies the NetBackup master server in that domain.
3	Target master (Domain 2)	NetBackup imports the image immediately, based on an SLP that contains an import operation. NetBackup can import the image quickly because the metadata is replicated as part of the image. (This import process is not the same as the import process available in the Catalog utility.)
4	Target master (Domain 2)	After the image is imported into the target domain, NetBackup continues to manage the copies in that domain. Depending on the configuration, the media server in Domain 2 can replicate the images to a media server in Domain 3.

About the domain relationship

For media server deduplication pools, the relationship between the originating domain and the target domain or domains is established by setting the properties in the source storage server.

Specifically, in the **Replication** tab of the **Change Storage Server** dialog box to configure the MSDP storage server.

See [“About the replication topology for Auto Image Replication”](#) on page 97.

Caution: Choose the target storage server or servers carefully. A target storage server must not also be a storage server for the originating domain.

One-to-many Auto Image Replication model

In this configuration, all copies are made in parallel. The copies are made within the context of one NetBackup job and simultaneously within the originating storage server context. If one target storage server fails, the entire job fails and is retried later.

All copies have the same **Target Retention**. To achieve different **Target Retention** settings in each target master server domain, either create multiple source copies or cascade duplication to target master servers.

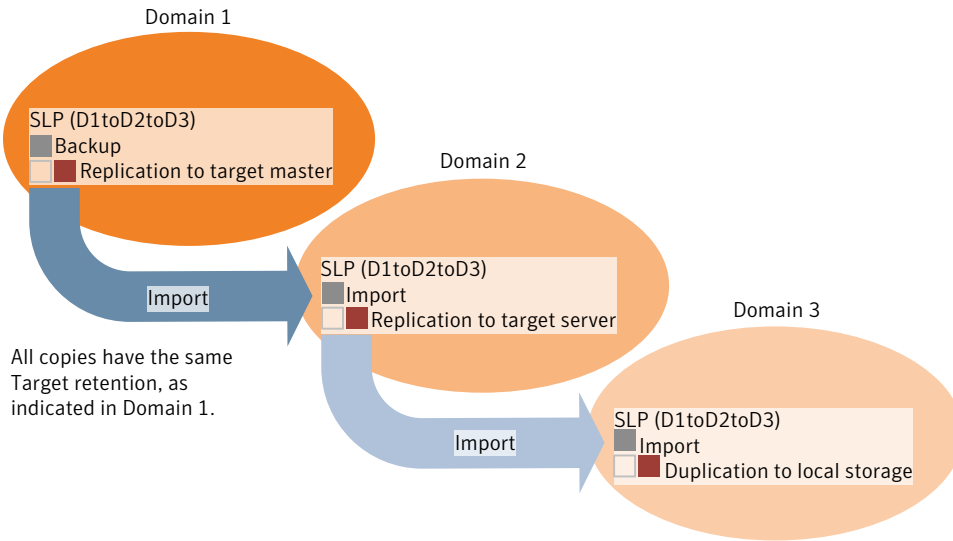
Cascading Auto Image Replication model

Replications can be cascaded from the originating domain to multiple domains. To do so, storage lifecycle policies are set up in each domain to anticipate the originating image, import it and then replicate it to the next target master.

[Figure 2-12](#) represents the following cascading configuration across three domains.

- The image is created in Domain 1, and then replicated to the target Domain 2.
- The image is imported in Domain 2, and then replicated to a target Domain 3.
- The image is then imported into Domain 3.

Figure 2-12 Cascading Auto Image Replication



In the cascading model, the originating master server for Domain 2 and Domain 3 is the master server in Domain 1.

Note: When the image is replicated in Domain 3, the replication notification event initially indicates that the master server in Domain 2 is the originating master server. However, when the image is successfully imported into Domain 3, this information is updated to correctly indicate that the originating master server is in Domain 1.

The cascading model presents a special case for the Import SLP that will replicate the imported copy to a target master. (This is the master server that is neither the first nor the last in the string of target master servers.)

As discussed previously, the requirements for an Import SLP include at least one operation that uses a **Fixed** retention type and at least one operation that uses a **Target Retention** type. So that the Import SLP can satisfy these requirements, the import operation must use a **Target Retention**.

Table 2-8 shows the difference in the import operation setup.

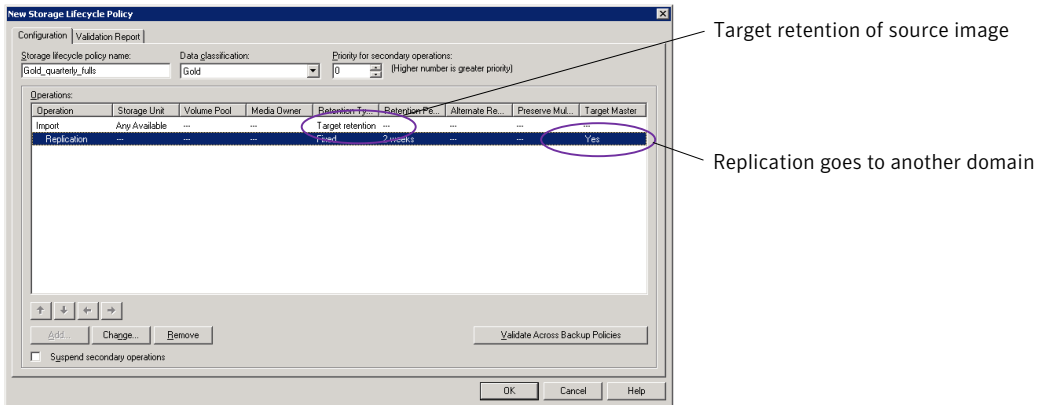
Table 2-8 Import operation difference in an SLP configured to replicate the imported copy

Import operation criteria	Import operation in a cascading model
The first operation must be an import operation.	Same; no difference.
A replication to target master must use a Fixed retention type	Same; no difference.
At least one operation must use the Target retention .	Here is the difference: To meet the criteria, the import operation must use Target retention .

The target retention is embedded in the source image.

Because the imported copy is the copy being replicated to a target master server domain, the fixed retention (three weeks in this example) on the replication to target master operation is ignored. The target retention is used instead. (See [Figure 2-13](#).)

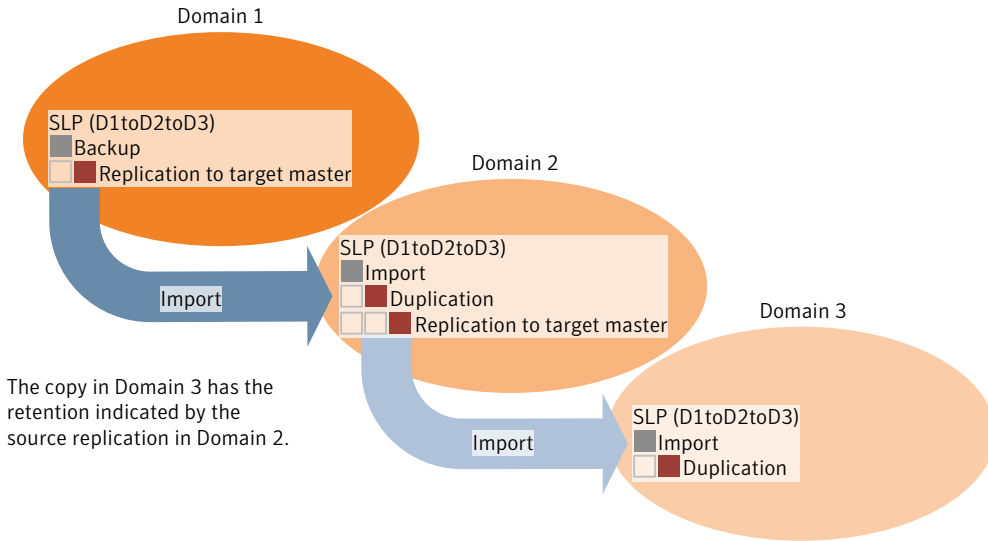
Figure 2-13 Storage lifecycle policy configured to replicate the imported copy



In the cascading model that is represented in [Figure 2-12](#), all copies have the same **Target Retention**—the **Target Retention** indicated in Domain 1.

For the copy in Domain 3 to have a different target retention, add an intermediary replication operation to the Domain 2 storage lifecycle policy. The intermediary replication operation acts as the source for the replication to target master. Since the target retention is embedded in the source image, the copy in Domain 3 honors the retention level that is set for the intermediary replication operation.

Figure 2-14 Cascading replications to target master servers, with various target retentions



About deduplication performance

Many factors affect performance, especially the server hardware and the network capacity.

[Table 2-9](#) provides information about performance during backup jobs for a deduplication storage server. The deduplication storage server conforms to the minimum host requirements. Client deduplication or load balancing servers are not used.

See [“About deduplication server requirements”](#) on page 26.

Table 2-9 Deduplication job load performance for a deduplication storage server

When	Description
Initial seeding	<p>Initial seeding is when all clients are first backed up.</p> <p>Approximately 15 to 20 jobs can run concurrently under the following conditions:</p> <ul style="list-style-type: none"> ■ The hardware meets minimum requirements. (More capable hardware improves performance.) ■ No compression. If data is compressed, the CPU usage increases quickly, which reduces the number of concurrent jobs that can be handled. ■ The deduplication rate is between 50% to 100%. The deduplication rate is the percentage of data already stored so it is not stored again. ■ The amount of data that is stored is less than 30% of the capacity of the storage.
Normal operation	<p>Normal operation is when all clients have been backed up once.</p> <p>Approximately 15 to 20 jobs can run concurrently and with high performance under the following conditions:</p> <ul style="list-style-type: none"> ■ The hardware meets minimum requirements. (More capable hardware improves performance.) ■ No compression. If data is compressed, the CPU usage increases quickly, which reduces the number of concurrent jobs that can be handled. ■ The deduplication rate is between 10% and 50%. The deduplication rate is the percentage of data already stored so it is not stored again. ■ The amount of data that is stored is between 30% to 90% of the capacity of the storage.
Clean up periods	<p>Clean up is when the NetBackup Deduplication Engine performs maintenance such as deleting expired backup image data segments.</p> <p>NetBackup maintains the same number of concurrent backup jobs as during normal operation. However, the average time to complete the jobs increases significantly.</p>
Storage approaches full capacity	<p>NetBackup maintains the same number of concurrent backup jobs as during normal operation under the following conditions:</p> <ul style="list-style-type: none"> ■ The hardware meets minimum requirements. (More capable hardware improves performance.) ■ The amount of data that is stored is between 85% to 90% of the capacity of the storage. <p>However, the average time to complete the jobs increases significantly.</p>

How file size may affect the deduplication rate

The small file sizes that are combined with large file segment sizes may result in low initial deduplication rates. However, after the deduplication engine performs file fingerprint processing, deduplication rates improve. For example, a second backup of a client shortly after the first does not show high deduplication rates. But the deduplication rate improves if the second backup occurs after the file fingerprint processing.

How long it takes the NetBackup Deduplication Engine to process the file fingerprints varies.

About deduplication stream handlers

NetBackup provides the stream handlers that process various backup data stream types. Stream handlers improve backup deduplication rates by processing the underlying data stream.

For data that has already been deduplicated, the first backup with a new stream handler produces a lower deduplication rate. After that first backup, the deduplication rate should surpass the rate from before the new stream handler was used.

Symantec continues to develop additional stream handlers to improve backup deduplication performance.

Deployment best practices

Symantec recommends that you consider the following practices when you implement NetBackup deduplication.

Use fully qualified domain names

Symantec recommends that you use fully qualified domain names for your NetBackup servers (and by extension, your deduplication servers). Fully qualified domain names can help to avoid host name resolution problems, especially if you use client-side deduplication.

Deduplication servers include the storage server and the load balancing servers (if any).

See [“Media write error \(84\)”](#) on page 200.

About scaling deduplication

You can scale deduplication processing to improve performance by using load balancing servers or client deduplication or both.

If you configure load balancing servers, those servers also perform deduplication. The deduplication storage server still functions as both a deduplication server and as a storage server. NetBackup uses standard load balancing criteria to select a load balancing server for each job. However, deduplication fingerprint calculations are not part of the load balancing criteria.

To completely remove the deduplication storage server from deduplication duties, do the following for every storage unit that uses the deduplication disk pool:

- Select **Only use the following media servers.**
- Select all of the load balancing servers but do not select the deduplication storage server.

The deduplication storage server performs storage server tasks only: storing and managing the deduplicated data, file deletion, and optimized duplication.

If you configure client deduplication, the clients deduplicate their own data. Some of the deduplication load is removed from the deduplication storage server and loading balancing servers.

Symantec recommends the following strategies to scale deduplication:

- For the initial full backups of your clients, use the deduplication storage server. For subsequent backups, use load balancing servers.
- Enable client-side deduplication gradually.
If a client cannot tolerate the deduplication processing workload, be prepared to move the deduplication processing back to a server.

Send initial full backups to the storage server

If you intend to use load balancing servers or client deduplication, use the storage server for the initial full backups of the clients. Then, send subsequent backups through the load balancing servers or use client deduplication for the backups.

Deduplication uses the same file fingerprint list regardless of which host performs the deduplication. So you can deduplicate data on the storage server first, and then subsequent backups by another host use the same fingerprint list. If the deduplication plug-in can identify the last full backup for the client and the policy combination, it retrieves the fingerprint list from the server. The list is placed in the fingerprint cache for the new backup.

See [“About deduplication fingerprinting”](#) on page 223.

Symantec also recommends that you implement load balancing servers and client deduplication gradually. Therefore, it may be beneficial to use the storage server for backups while you implement deduplication on other hosts.

Increase the number of jobs gradually

Symantec recommends that you increase the **Maximum concurrent jobs** value gradually. (The **Maximum concurrent jobs** is a storage unit setting.) The initial backup jobs (also known as initial seeding) require more CPU and memory than successive jobs. After initial seeding, the storage server can process more jobs concurrently. Gradually increase the jobs value over time.

Testing shows that the upper limit for a storage server with 8 GB of memory and 4 GB of swap space is 50 concurrent jobs.

Introduce load balancing servers gradually

Symantec recommends that you add load balancing servers only after the storage server reaches maximum CPU utilization. Then, introduce load balancing servers one at a time. It may be easier to evaluate how your environment handles traffic and easier to troubleshoot any problems with fewer hosts added for deduplication.

Many factors affect deduplication server performance.

Because of the various factors, Symantec recommends that you maintain realistic expectations about using multiple servers for deduplication. If you add one media server as a load balancing server, overall throughput should be faster. However, adding one load balancing server may not double the overall throughput rate, adding two load balancing servers may not triple the throughput rate, and so on.

If all of the following apply to your deduplication environment, your environment may be a good candidate for load balancing servers:

- The deduplication storage server is CPU limited on any core.
- Memory resources are available on the storage server.
- Network bandwidth is available on the storage server.
- Back-end I/O bandwidth to the deduplication pool is available.
- Other NetBackup media servers have CPU available for deduplication.

Gigabit Ethernet should provide sufficient performance in many environments. If your performance objective is the fastest throughput possible with load balancing servers, you should consider 10 Gigabit Ethernet.

Implement client deduplication gradually

If you configure clients to deduplicate their own data, do not enable all of those clients at the same time. Implement client deduplication gradually, as follows:

- Use the storage server for the initial backup of the clients.
- Enable deduplication on only a few clients at a time.
It may be easier to evaluate how your environment handles traffic and easier to troubleshoot any problems with fewer hosts added for deduplication.

If a client cannot tolerate the deduplication processing workload, be prepared to move the deduplication processing back to the storage server.

Use deduplication compression and encryption

Do not use compression or encryption in a NetBackup policy; rather, use the compression or the encryption that is part of the deduplication process.

See “[About deduplication compression](#)” on page 34.

About the optimal number of backup streams

A backup stream appears as a separate job in the NetBackup Activity Monitor. Various methods exist to produce streams. In NetBackup, you can use backup policy settings to configure multiple streams. The NetBackup for Oracle agent lets you configure multiple streams; also for Oracle the RMAN utilities can provide multiple backup channels.

For client deduplication, the optimal number of backup streams is two.

Media server deduplication can process multiple streams on multiple cores simultaneously. For large datasets in applications such as Oracle, media server deduplication leverages multiple cores and multiple streams. Therefore, media server deduplication may be a better solution when the application can provide multiple streams or channels.

More detailed information about backup streams is available.

<http://www.symantec.com/docs/TECH77575>

About storage unit groups for deduplication

You can use a storage unit group as a backup destination for NetBackup deduplication, as follows:

- A storage unit group can contain the storage units that have a **Media Server Deduplication Pool** as the storage destination.

- A storage unit group can contain the storage units that have a PureDisk storage pool as the storage destination.

A group must contain storage units of one storage destination type only. That is, a group cannot contain both **Media Server Deduplication Pool** storage units and PureDisk storage pool storage units.

Storage unit groups avoid a single point of failure that can interrupt backup service.

The best storage savings occur when a backup policy stores its data in the same deduplication destination disk pool instead of across multiple disk pools. For this reason, the **Failover** method for the **Storage unit selection** uses the least amount of storage. All of the other methods are designed to use different storage every time the backup runs.

Note: NetBackup does not support storage unit groups for optimized duplication of deduplicated data. If you use a storage unit group as a destination for optimized duplication of deduplicated data, NetBackup uses regular duplication.

About protecting the deduplicated data

Symantec recommends the following methods to protect the deduplicated backup data and the deduplication database:

- Use NetBackup optimized duplication to copy the images to another deduplication node off-site location.

Optimized duplication copies the primary backup data to another deduplication pool. It provides the easiest, most efficient method to copy data off-site yet remain in the same NetBackup domain. You then can recover from a disaster that destroys the storage on which the primary copies reside by retrieving images from the other deduplication pool.

See [“About MSDP optimized duplication within the same domain”](#) on page 36.

See [“Configuring optimized duplication of deduplicated data”](#) on page 94.

- For the primary deduplication storage, use a SAN volume with resilient storage methodologies to replicate the data to a remote site. If the deduplication database is on a different SAN volume, replicate that volume to the remote site also.

The preceding methods help to make your data highly-available.

Also, you can use NetBackup to back up the deduplication storage server system or program disks. If the disk on which NetBackup resides fails and you have to replace it, you can use NetBackup to restore the media server.

Save the deduplication storage server configuration

Symantec recommends that you save the storage server configuration. Getting and saving the configuration can help you with recovery of your environment. For disaster recovery, you may need to set the storage server configuration by using a saved configuration file.

If you save the storage server configuration, you must edit it so that it includes only the information that is required for recovery.

See [“About saving the deduplication storage server configuration”](#) on page 128.

See [“Saving the deduplication storage server configuration”](#) on page 129.

See [“Editing a deduplication storage server configuration file”](#) on page 129.

Plan for disk write caching

Storage components may use hardware caches to improve read and write performance. Among the storage components that may use a cache are disk arrays, RAID controllers, or the hard disk drives themselves.

If your storage components use caches for disk write operations, ensure that the caches are protected from power fluctuations or power failure. If you do not protect against power fluctuations or failure, data corruption or data loss may occur.

Protection can include the following:

- A battery backup unit that supplies power to the cache memory so write operations can continue if power is restored within sufficient time.
- An uninterruptible power supply that allows the components to complete their write operations.

If your devices that have caches are not protected, Symantec recommends that you disable the hardware caches. Read and write performance may decline, but you help to avoid data loss.

How deduplication restores work

The data is first reassembled on a media server before it is restored, even for the clients that deduplicate their own data. The media server that performs the restore always is a deduplication server (that is, hosts the deduplication plug-in).

The backup server may not be the server that performs the restore. If another server has credentials for the NetBackup Deduplication Engine (that is, for the storage server), NetBackup may use that server for the restore job. NetBackup chooses the least busy server for the restore.

The following other servers can have credentials for the NetBackup Deduplication Engine:

- A load balancing server in the same deduplication node.
- A deduplication server in a different deduplication node that is the target of optimized duplication.

Optimized duplication requires a server in common between the two deduplication nodes.

See [“About the media servers for optimized MSDP duplication within the same domain”](#) on page 38.

You can specify the server to use for restores.

See [“Specifying the restore server”](#) on page 184.

Replacing the PureDisk Deduplication Option with Media Server Deduplication on the same host

The PureDisk Deduplication Option provides deduplication of NetBackup backups for NetBackup release 6.5. The destination storage for PDDO is a PureDisk storage pool. The PureDisk agent that performs the deduplication is installed from the PureDisk software distribution, not from the NetBackup distribution. PDDO is not the same as integrated NetBackup deduplication.

You can upgrade to 7.0 a NetBackup media server that hosts a PDDO agent and use that server for integrated NetBackup deduplication. The storage can remain the PureDisk storage pool, and NetBackup maintains access to all of the valid backup images in the PureDisk storage pool.

If you perform this procedure, the NetBackup deduplication plug-in replaces the PureDisk agent on the media server. The NetBackup deduplication plug-in can deduplicate data for either integrated NetBackup deduplication or for a PureDisk storage pool. The PDDO agent can deduplicate data only for a PureDisk storage pool.

Note: To use the NetBackup deduplication plug-in with a PureDisk storage pool, the PureDisk storage pool must be part of a PureDisk 6.6 or later environment.

Table 2-10 Replacing a PDDO host with a media server deduplication host

Step	Task	Procedure
Step 1	Deactivate all backup policies that use the host	Deactivate the policies to ensure that no activity occurs on the host. See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I.</i>
Step 2	Remove the PDDO plug-in	NetBackup deduplication components cannot reside on the same host as a PureDisk Deduplication Option (PDDO) agent. Therefore, remove the PDDO agent from the host. See the <i>NetBackup PureDisk Deduplication Option Guide</i> .
Step 3	Upgrade the media server to 7.0 or later	If the media server runs a version of NetBackup earlier than 7.0, upgrade that server to NetBackup 7.0 or later. See the <i>NetBackup Installation Guide for UNIX and Linux</i> . See the <i>NetBackup Installation Guide for Windows</i> .
Step 4	Configure the host	In the Storage Server Configuration Wizard , select PureDisk Deduplication Pool and enter the name of the Storage Pool Authority. See “Configuring a NetBackup deduplication storage server” on page 79.
Step 5	Activate your backup policies	See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I.</i>

Migrating from PureDisk to the NetBackup Media Server Deduplication option

NetBackup cannot use the storage hardware while PureDisk uses it for storage. The structure of the PureDisk storage is different than the structure of the storage for integrated NetBackup deduplication. The disk systems cannot be used simultaneously by both NetBackup and PureDisk. The PureDisk images on the storage cannot be transferred to the deduplication storage server storage.

Therefore, to migrate from NetBackup PureDisk to the NetBackup Media Server Deduplication Option, Symantec recommends that you age the PureDisk storage pool backups until they expire.

Table 2-11 To migrate from PureDisk to NetBackup deduplication

Step	Task	Procedure
Step 1	Install and configure NetBackup	See the <i>NetBackup Installation Guide for UNIX and Linux</i> . See the <i>NetBackup Installation Guide for Windows</i> .
Step 2	Configure NetBackup deduplication	See “ Configuring NetBackup media server deduplication ” on page 76.
Step 3	Redirect your backup jobs	Redirect your backup jobs to the NetBackup media server deduplication pool. See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> . See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .
Step 4	Uninstall PureDisk	After the PureDisk backup images expire, uninstall PureDisk. See your NetBackup PureDisk documentation.

See “[Migrating from another storage type to deduplication](#)” on page 62.

Migrating from another storage type to deduplication

To migrate from another NetBackup storage type to deduplication storage, Symantec recommends that you age the backup images on the other storage until they expire. Symantec recommends that you age the backup images if you migrate from disk storage or tape storage.

You should not use the same disk storage for NetBackup deduplication while you use it for other storage such as AdvancedDisk, BasicDisk, or SharedDisk. Each type manages the storage differently and each requires exclusive use of the storage. Also, the NetBackup Deduplication Engine cannot read the backup images that another NetBackup storage type created. Therefore, you should age the data so it expires before you repurpose the storage hardware. Until that data expires, two storage destinations exist: the media server deduplication pool and the other storage. After the images on the other storage expire and are deleted, you can repurpose it for other storage needs.

Table 2-12 Migrating to NetBackup deduplication

Step	Task	Procedure
Step 1	Configure NetBackup deduplication	See “Configuring NetBackup media server deduplication” on page 76.
Step 2	Redirect your backup jobs	<p>Redirect your backup jobs to the media server deduplication pool storage unit. To do so, change the backup policy storage destination to the storage unit for the deduplication pool.</p> <p>See the <i>NetBackup Administrator’s Guide for UNIX and Linux, Volume I</i>.</p> <p>See the <i>NetBackup Administrator’s Guide for Windows, Volume I</i>.</p>
Step 3	Repurpose the storage	<p>After all of the backup images that are associated with the storage expire, repurpose that storage.</p> <p>If it is disk storage, you cannot add it to an existing media server deduplication pool. You can use it as storage for another, new deduplication node.</p>

See [“Migrating from PureDisk to the NetBackup Media Server Deduplication option”](#) on page 61.

Provisioning the storage

This chapter includes the following topics:

- [About provisioning the deduplication storage](#)
- [About deduplication storage requirements](#)
- [About deduplication storage capacity](#)
- [About the deduplication storage paths](#)
- [Do not modify storage directories and files](#)
- [About adding additional storage](#)
- [About volume management for NetBackup deduplication](#)

About provisioning the deduplication storage

How to provision the storage is beyond the scope of the NetBackup documentation. For help, consult the storage vendor's documentation.

What you choose as your storage destination affects how you provision the storage. NetBackup requirements also may affect how you provision the storage.

How many storage instances you provision depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication, as follows:

Optimized duplication within the same NetBackup domain You must provision the storage for at least two deduplication nodes in the same NetBackup domain:

- Storage for the backups, which is the source for the duplication operations.
- Different storage in another deduplication node for the copies of the backup images, which is the target for the duplication operations.

Auto Image Replication to a different NetBackup domain You must provision the storage in at least two NetBackup domains:

- Storage for the backups in the originating domain. This storage contains your client backups. It is the source for the duplication operations.
- Different storage in the remote domain for the copies of the backup images. This storage is the target for the replication operations that run in the originating domain.

See [“About NetBackup Auto Image Replication”](#) on page 47.

See [“About the deduplication storage destination”](#) on page 22.

See [“Planning your deduplication deployment”](#) on page 20.

About deduplication storage requirements

The following describes the storage requirements for the NetBackup Media Server Deduplication Option:

Table 3-1 Deduplication storage requirements

Component	Requirements
Storage media	<p>Disk, with the following minimum requirements per individual data stream (read or write):</p> <ul style="list-style-type: none"> ■ Up to 32 TBs of storage: <ul style="list-style-type: none"> ■ 130 MB/sec. ■ 200 MB/sec for enterprise-level performance. ■ 32 to 48 TBs of storage: <ul style="list-style-type: none"> 200 MB/sec. <p>Symantec recommends that you store the data and the deduplication database on separate disk, each with 200 MB/sec read or write speed.</p> ■ 48 to 64 TBs of storage: <ul style="list-style-type: none"> 250 MB/sec. <p>Symantec recommends that you store the data and the deduplication database on separate disk, each with 250 MB/sec read or write speed.</p> <p>These are minimum requirements for single stream read or write performance. Greater individual data stream capability or aggregate capability may be required to satisfy your objectives for writing to and reading from disk.</p>
Connection	<p>Storage area network (Fibre Channel or iSCSI), direct-attached storage (DAS), or internal disks.</p> <p>The storage area network should conform to the following criteria:</p> <ul style="list-style-type: none"> ■ A dedicated, low latency network for storage with a maximum 0.1-millisecond latency per round trip. ■ Enough bandwidth on the storage network to satisfy your throughput objectives. Symantec supports iSCSI on storage networks with at least 10-Gigabit Ethernet network bandwidth. Symantec recommends Fibre Channel storage networks with at least 4-Gigabit network bandwidth. ■ The storage server should have an HBA or HBAs dedicated to the storage. Those HBAs must have enough bandwidth to satisfy your throughput objectives. <p>Local disk storage may leave you vulnerable in a disaster. SAN disk can be remounted at a newly provisioned server with the same name.</p>

NetBackup requires exclusive use of the disk resources. If the storage is used for purposes other than backups, NetBackup cannot manage disk pool capacity or manage storage lifecycle policies correctly. Therefore, NetBackup must be the only entity that uses the storage.

NetBackup deduplication does not support file based storage protocols, such as CIFS or NFS, for deduplication storage.

The storage must be configured and operational before you can configure deduplication in NetBackup.

About support for more than 32-TB of storage

Beginning with the 7.5 release, NetBackup supports up to 64 TB of storage space.

To use more storage space may require that you upgrade the storage server host CPU and memory.

See [“About deduplication server requirements”](#) on page 26.

If you upgrade an existing NetBackup deduplication environment, do the following to use up to 64 TB of storage space:

- Stop the deduplication services (`spad` and `spoold`) on the storage server host.
- Grow the disk storage to the desired size up to 64 TBs.
See [“About adding additional storage”](#) on page 70.
- If more CPU and memory are required, shut down the storage server host and then upgrade the CPU and add the required amount of memory.
- Start the host.
- Upgrade NetBackup on the storage server host.

About deduplication storage capacity

The maximum deduplication storage capacity is 64 TBs. NetBackup reserves 4 percent of the storage space for the deduplication database and transaction logs. Therefore, a storage full condition is triggered at a 96 percent threshold.

For performance optimization, Symantec recommends that you use a separate disk, volume, partition, or spindle for the deduplication database. If you use separate storage for the deduplication database, NetBackup still uses the 96 percent threshold to protect the data storage from any possible overload.

If your storage requirements exceed the capacity of a media server deduplication node, do one of the following:

- Use more than one media server deduplication node.
- Use a PureDisk storage pool as the storage destination.
A PureDisk storage pool provides larger storage capacity; PureDisk also provides global deduplication.
See [“About the deduplication storage destination”](#) on page 22.

Only one deduplication storage path can exist on a media server. You cannot add another storage path to increase capacity beyond 64 TBs.

About the deduplication storage paths

When you configure the deduplication storage server, you must enter the path name to the storage. The storage path is the directory in which NetBackup stores the raw backup data.

Because the storage requires a directory path, do not use only a root node (/) or drive letter (C:\) as the storage path.

You also can specify a different location for the deduplication database. The database path is the directory in which NetBackup stores and maintains the structure of the stored deduplicated data.

For performance optimization, Symantec recommends that you use a separate disk, volume, partition, or spindle for the deduplication database. Depending on the size of your deduplication storage, Symantec also recommends a separate path for the deduplication database.

See [“About deduplication storage requirements”](#) on page 66.

If the directory or directories do not exist, NetBackup creates them and populates them with the necessary subdirectory structure. If the directory or directories exist, NetBackup populates them with the necessary subdirectory structure.

The path names must use ASCII characters only.

Caution: You cannot change the paths after NetBackup configures the deduplication storage server. Therefore, carefully decide during the planning phase where and how you want the deduplicated backup data stored.

Do not modify storage directories and files

Unless you are directed to do so by the NetBackup documentation or by a Symantec support representative, do not do the following:

- Add files to the deduplication storage directories or database directories.
- Delete files from the deduplication storage directories or database directories.
- Modify files in the deduplication storage directories or database directories.
- Move files within the deduplication storage directories or database directories.

Failure to follow these directives can result in data loss.

About adding additional storage

The storage for a NetBackup **Media Server Deduplication Pool** is exposed as a single disk volume. You cannot add another volume to an existing **Media Server Deduplication Pool**.

To increase the capacity of a **Media Server Deduplication Pool**, grow the existing volume.

See “[Resizing the deduplication storage partition](#)” on page 182.

About volume management for NetBackup deduplication

If you use a tool to manage the volumes for NetBackup **Media Server Deduplication Pool** storage, Symantec recommends that you use the Veritas Storage Foundation. Storage Foundation includes the Veritas Volume Manager and the Veritas File System.

For supported systems, see the Storage Foundation hardware compatibility list at the Symantec Web site:

<http://www.symantec.com/>

Note: Although Storage Foundation supports NFS, NetBackup does not support NFS targets for **Media Server Deduplication Pool** storage. Therefore, **Media Server Deduplication Pool** does not support NFS with Storage Foundation.

Licensing deduplication

This chapter includes the following topics:

- [About licensing deduplication](#)
- [About the deduplication license key](#)
- [Licensing NetBackup deduplication](#)

About licensing deduplication

The NetBackup deduplication components are installed by default on the supported host systems. However, you must enter a license key to enable deduplication.

Before you try to install or upgrade to a NetBackup version that supports deduplication, be aware of the following:

- NetBackup supports deduplication on specific 64-bit host operating systems. If you intend to upgrade an existing media server and use it for deduplication, that host must be supported.

For the supported systems, see the *NetBackup Release Notes*.

- NetBackup deduplication components cannot reside on the same host as a PureDisk Deduplication Option agent.

To use a PDDO agent host for NetBackup deduplication, first remove the PDDO agent from that host.

See the *NetBackup PureDisk Deduplication Option (PDDO) Guide*.

Then, upgrade that host to NetBackup 7.0 or later.

Finally, configure that host as a deduplication storage server or as a load balancing server.

About the deduplication license key

NetBackup deduplication is licensed separately from base NetBackup.

The NetBackup Deduplication Option license key enables both NetBackup media server deduplication and NetBackup client deduplication. The license is a front-end capacity license. It is based on the size of the data to be backed up, not on the size of the deduplicated data.

You may have a single license key that activates both NetBackup and optional features. Alternatively, you may have one license key that activates NetBackup and another key that activates deduplication.

If you remove the NetBackup Deduplication Option license key or if it expires, you cannot create new deduplication disk pools. you also cannot create the storage units that reference NetBackup deduplication pools.

NetBackup does not delete the disk pools or the storage units that reference the disk pools. You can use them again if you enter a valid license key.

The NetBackup Deduplication Option license key also enables the **Use accelerator** feature on the NetBackup policy **Attributes** tab. Accelerator increases the speed of full backups for files systems. Accelerator works with deduplication storage units as well as with other storage units that do not require the deduplication option. More information on Accelerator is available.

See the *NetBackup Administrator's Guide for UNIX and Linux, Volume I*.

See the *NetBackup Administrator's Guide for Windows, Volume I*.

Licensing NetBackup deduplication

If you installed the license key for deduplication when you installed or upgraded NetBackup, you do not need to perform this procedure.

Enter the license key on the NetBackup master server. The following procedure describes how to use the NetBackup Administration Console to enter the license key.

To license NetBackup deduplication

- 1 On the **Help** menu of the **NetBackup Administration Console**, select **License Keys**.
- 2 In the **NetBackup License Keys** dialog box, click **New**.
- 3 In the **Add a New License Key** dialog box, enter the license key and click **Add** or **OK**.

- 4 Click **Close**.
- 5 Restart all the NetBackup services and daemons.

Configuring deduplication

This chapter includes the following topics:

- [Configuring NetBackup media server deduplication](#)
- [Configuring NetBackup client-side deduplication](#)
- [Configuring a NetBackup deduplication storage server](#)
- [About NetBackup deduplication pools](#)
- [Configuring a deduplication disk pool](#)
- [Configuring a deduplication storage unit](#)
- [Enabling deduplication encryption](#)
- [Configuring optimized synthetic backups for deduplication](#)
- [About configuring optimized duplication and replication bandwidth](#)
- [Configuring MSDP optimized duplication copy behavior](#)
- [Configuring a separate network path for MSDP optimized duplication](#)
- [Configuring optimized duplication of deduplicated data](#)
- [About the replication topology for Auto Image Replication](#)
- [Configuring a target for MSDP replication](#)
- [Viewing the replication topology for Auto Image Replication](#)
- [About the storage lifecycle policies required for Auto Image Replication](#)
- [Creating a storage lifecycle policy](#)
- [About backup policy configuration](#)

- [Creating a policy using the Policy Configuration Wizard](#)
- [Creating a policy without using the Policy Configuration Wizard](#)
- [Enabling client-side deduplication](#)
- [Resilient Network properties](#)
- [Specifying resilient connections](#)
- [Seeding the fingerprint cache for remote client-side deduplication](#)
- [Adding a deduplication load balancing server](#)
- [About the pd.conf configuration file for NetBackup deduplication](#)
- [Editing the pd.conf deduplication file](#)
- [About the contentrouter.cfg file for NetBackup deduplication](#)
- [About saving the deduplication storage server configuration](#)
- [Saving the deduplication storage server configuration](#)
- [Editing a deduplication storage server configuration file](#)
- [Setting the deduplication storage server configuration](#)
- [About the deduplication host configuration file](#)
- [Deleting a deduplication host configuration file](#)
- [Resetting the deduplication registry](#)
- [Configuring deduplication log file timestamps on Windows](#)
- [Setting NetBackup configuration options by using bpsetconfig](#)

Configuring NetBackup media server deduplication

This topic describes how to configure media server deduplication in NetBackup.

[Table 5-1](#) describes the configuration tasks.

The NetBackup administrator's guides describe how to configure a base NetBackup environment.

See the *NetBackup Administrator's Guide for Windows, Volume I*.

See the *NetBackup Administrator's Guide for UNIX and Linux, Volume I*.

Table 5-1 Media server deduplication configuration tasks

Step	Task	Procedure
Step 1	Install the license key	See “Licensing NetBackup deduplication” on page 72.
Step 2	Configure a deduplication storage server	How many storage servers you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication. See “About NetBackup deduplication servers” on page 25. See “Configuring a NetBackup deduplication storage server” on page 79.
Step 3	Configure a disk pool	How many disk pools you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication. See “About NetBackup deduplication pools” on page 79. See “Configuring a deduplication disk pool” on page 81.
Step 4	Configure a storage unit	See “Configuring a deduplication storage unit” on page 83.
Step 5	Enable encryption	Encryption is optional. See “Enabling deduplication encryption” on page 87.
Step 6	Configure optimized synthetic backups	Optimized synthetic backups are optional. See “Configuring optimized synthetic backups for deduplication” on page 89.
Step 7	Configure optimized duplication copy	Optimized duplication is optional. See “Configuring a separate network path for MSDP optimized duplication” on page 92. See “Configuring optimized duplication of deduplicated data” on page 94.
Step 8	Configure replication	Replication is optional. See “About the replication topology for Auto Image Replication” on page 97. See “Configuring a target for MSDP replication” on page 98. See “Viewing the replication topology for Auto Image Replication” on page 100. See “About the storage lifecycle policies required for Auto Image Replication” on page 102. See “Creating a storage lifecycle policy” on page 105.

Table 5-1 Media server deduplication configuration tasks (*continued*)

Step	Task	Procedure
Step 9	Configure a backup policy	Use the deduplication storage unit as the destination for the backup policy. See “Creating a policy using the Policy Configuration Wizard” on page 111. See “Creating a policy without using the Policy Configuration Wizard” on page 111.
Step 10	Specify advanced deduplication settings	Advanced settings are optional. See “About the pd.conf configuration file for NetBackup deduplication” on page 119. See “Editing the pd.conf deduplication file” on page 126. See “pd.conf file settings for NetBackup deduplication” on page 119.

Configuring NetBackup client-side deduplication

This topic describes how to configure client deduplication in NetBackup.

Table 5-2 Client deduplication configuration tasks

Step	Task	Procedure
Step 1	Configure media server deduplication	See “Configuring NetBackup media server deduplication” on page 76.
Step 2	Learn about client deduplication	See “About NetBackup Client Deduplication” on page 28.
Step 3	Configure a resilient connection for remote office clients	Resilient connections are optional. See “About remote office client deduplication” on page 31. See “Resilient Network properties” on page 113. See “Specifying resilient connections” on page 116.
Step 4	Enable client-side deduplication	See “Enabling client-side deduplication” on page 112.
Step 5	Seed the fingerprint cache of a remote client-side deduplication client	Remote client seeding is optional. See “Seeding the fingerprint cache for remote client-side deduplication” on page 117.

Configuring a NetBackup deduplication storage server

Configure in this context means to configure a NetBackup media server as a storage server for deduplication.

See [“About NetBackup deduplication servers”](#) on page 25.

When you configure a storage server for deduplication, you specify the following:

- The type of storage server.
For NetBackup media server deduplication, select **Media Server Deduplication Pool** for the type of disk storage.
For a PureDisk deduplication pool, select **PureDisk Deduplication Pool** for the type of disk storage.
- The credentials for the deduplication engine.
See [“About NetBackup Deduplication Engine credentials”](#) on page 32.
- The storage paths.
See [“About the deduplication storage paths”](#) on page 69.
- The network interface.
See [“About the network interface for deduplication”](#) on page 33.
- The load balancing servers, if any.
See [“About NetBackup deduplication servers”](#) on page 25.

When you create the storage server, the wizard lets you create a disk pool and storage unit also.

To configure a deduplication storage server by using the wizard

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Configure Disk Storage Servers**.
- 2 Follow the wizard screens to configure a deduplication storage server.
- 3 After NetBackup creates the deduplication storage server, you can click **Next** to continue to the **Disk Pool Configuration Wizard**.

About NetBackup deduplication pools

Deduplication pools are the disk pools that are the storage destination for deduplicated backup data. NetBackup servers or NetBackup clients deduplicate the backup data that is stored in a deduplication pool.

Two types of deduplication pools exist, as follows:

- A NetBackup **Media Server Deduplication Pool** represents the disk storage that is attached to a NetBackup media server. NetBackup deduplicates the data and hosts the storage.
- A NetBackup **PureDisk Deduplication Pool** represents a PureDisk storage pool. NetBackup deduplicates the data, and PureDisk hosts the storage. A **PureDisk Deduplication Pool** destination requires that PureDisk be at release 6.6 or later.

When you configure a deduplication pool, choose **PureDisk** as the deduplication pool type.

NetBackup requires exclusive ownership of the disk resources that comprise the deduplication pool. If you share those resources with other users, NetBackup cannot manage deduplication pool capacity or storage lifecycle policies correctly.

How many deduplication pools you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication, as described in the following table:

Table 5-3 Deduplication pools for duplication or replication

Type	Requirements
Optimized duplication within the same NetBackup domain	<p>Optimized duplication in the same domain requires the following deduplication pools:</p> <ul style="list-style-type: none"> ■ At least one for the backup storage, which is the source for the duplication operations. The source deduplication pool is in one deduplication node. ■ Another to store the copies of the backup images, which is the target for the duplication operations. The target deduplication pool is in a different deduplication node.
Auto Image Replication to a different NetBackup domain	<p>Auto Image Replication deduplication pools can be either replication source or replication target. The replication properties denote the purpose of the deduplication pool. The deduplication pools inherit the replication properties from their volumes. See “About the replication topology for Auto Image Replication” on page 97.</p> <p>Auto Image Replication requires the following deduplication pools:</p> <ul style="list-style-type: none"> ■ At least one replication source deduplication pool in the originating domain. A replication source deduplication pool is one to which you send your backups. The backup images on the source deduplication pool are replicated to a deduplication pool in the remote domain or domains. ■ At least one replication target deduplication pool in a remote domain or domains. A replication target deduplication pool is the target for the duplication operations that run in the originating domain. <p>See “About NetBackup Auto Image Replication” on page 47.</p>

Configuring a deduplication disk pool

Before you can configure a NetBackup disk pool, a NetBackup deduplication storage server must exist.

See [“Configuring a NetBackup deduplication storage server”](#) on page 79.

The **Storage Server Configuration Wizard** lets you configure one disk pool. To configure additional disk pools, launch the **Disk Pool Configuration Wizard**.

When you configure a disk pool for deduplication, you specify the following:

- The type of disk pool (**PureDisk**).
- The NetBackup deduplication storage server to query for the disk storage to use for the pool.
- The disk volume to include in the pool.
NetBackup exposes the storage as a single volume.
- The disk pool properties.
See [“Media server deduplication pool properties”](#) on page 81.

Symantec recommends that disk pool names be unique across your enterprise.

To create a NetBackup disk pool

- 1 In the **NetBackup Administration Console**, select the **Media and Device Management** node.
- 2 From the list of wizards in the **Details** pane, click **Configure Disk Pool** and follow the wizard instructions.
For help, see the wizard help.
- 3 After NetBackup creates the deduplication pool, you have the option to create a storage unit that uses the pool.

Media server deduplication pool properties

[Table 5-4](#) describes the disk pool properties.

Table 5-4 Media server deduplication pool properties

Property	Description
Name	The disk pool name.
Storage server	The storage server name. The storage server is the same as the NetBackup media server to which the storage is attached.

Table 5-4 Media server deduplication pool properties (*continued*)

Property	Description
Disk volume	<p>For a media server deduplication pool, all disk storage is exposed as a single volume.</p> <p>PureDiskVolume is a virtual name for the storage that is contained within the directories you specified for the storage path and the database path.</p>
Refresh	<p>Query the storage server for its replication capabilities. If the capabilities are different than the current NetBackup configuration, update the configuration.</p>
Available space	<p>The amount of space available in the disk pool.</p>
Raw size	<p>The total raw size of the storage in the disk pool.</p>
Comment	<p>A comment that is associated with the disk pool.</p>
High water mark	<p>The High water mark setting is a threshold that invokes the following actions:</p> <ul style="list-style-type: none"> ■ The High water mark indicates that the PureDiskVolume is full. When the PureDiskVolume reaches the High water mark, NetBackup fails any backup jobs that are assigned to the storage unit. NetBackup also does not assign new jobs to a storage unit in which the disk pool is full. NetBackup also fails backup jobs if the PureDiskVolume does not contain enough storage for its estimated space requirement. ■ NetBackup begins image cleanup when the PureDiskVolume reaches the High water mark; image cleanup expires the images that are no longer valid. NetBackup again assigns jobs to the storage unit when image cleanup reduces the PureDiskVolume capacity to less than the High water mark. <p>The default is 98%.</p>
Low water mark	<p>The Low water mark has no affect on the PureDiskVolume.</p>

Table 5-4 Media server deduplication pool properties (*continued*)

Property	Description
Limit I/O streams	<p>Select to limit the number of read and write streams (that is, jobs) for each volume in the disk pool. A job may read backup images or write backup images. By default, there is no limit. If you select this property, also configure the number of streams to allow per volume.</p> <p>When the limit is reached, NetBackup chooses another volume for write operations, if available. If not available, NetBackup queues jobs until a volume is available.</p> <p>Too many streams may degrade performance because of disk thrashing. Disk thrashing is excessive swapping of data between RAM and a hard disk drive. Fewer streams can improve throughput, which may increase the number of jobs that complete in a specific time period.</p>
per volume	<p>Select or enter the number of read and write streams to allow per volume.</p> <p>Many factors affect the optimal number of streams. Factors include but are not limited to disk speed, CPU speed, and the amount of memory.</p>

Configuring a deduplication storage unit

Create one or more storage units that reference the disk pool.

A storage unit inherits the properties of the disk pool. If the storage unit inherits replication properties, the properties signal to a NetBackup storage lifecycle policy the intended purpose of the storage unit and the disk pool. Auto Image Replication requires storage lifecycle policies.

The **Disk Pool Configuration Wizard** lets you create a storage unit; therefore, you may have created a storage unit when you created a disk pool. To determine if storage units exist for the disk pool, see the **NetBackup Management > Storage > Storage Units** window of the Administration Console.

To configure a storage unit from the Actions menu

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Storage > Storage Units**.
- 2 On the **Actions** menu, select **New > Storage Unit**.

New Storage Unit

Storage unit name:
|

Storage unit type:
Disk On demand only

Disk type:
PureDisk

Storage unit properties

Disk Pool:
dp-msdp-source-domain-backups

Media Server:
 Use any available media server
 Only use the following media servers

Media Servers

- bit1.symantecs.org
- bit2.symantecs.org
- bit3.symantecs.org

Maximum concurrent jobs: 1
Maximum fragment size: 51200 Megabytes

- 3 Complete the fields in the **New Storage Unit** dialog box.

For a storage unit for optimized duplication destination, select **Only use the following media servers**. Then select the media servers that are common between the two deduplication nodes.

See [“Deduplication storage unit properties”](#) on page 85.

See “[Deduplication storage unit recommendations](#)” on page 86.

Deduplication storage unit properties

The following are the configuration options for a PureDisk disk pool storage unit.

Table 5-5 Deduplication storage unit properties

Property	Description
Storage unit name	A unique name for the new storage unit. The name can describe the type of storage. The storage unit name is the name used to specify a storage unit for policies and schedules. The storage unit name cannot be changed after creation.
Storage unit type	Select Disk as the storage unit type.
Disk type	Select PureDisk for the disk type for a media server deduplication pool, a PureDisk deduplication pool, or a PureDisk Deduplication Option storage pool.
Disk pool	Select the disk pool that contains the storage for this storage unit. All disk pools of the specified Disk type appear in the Disk pool list. If no disk pools are configured, no disk pools appear in the list.
Media server	The Media server setting specifies the NetBackup media servers that can deduplicate the data for this storage unit. Only the load balancing servers appear in the media server list. Specify the media server or servers as follows: <ul style="list-style-type: none"> ■ To allow any server in the media server list to deduplicate data, select Use any available media server. ■ To use specific media servers to deduplicate the data, select Only use the following media servers. Then, select the media servers to allow. NetBackup selects the media server to use when the policy runs.
Maximum fragment size	For normal backups, NetBackup breaks each backup image into fragments so it does not exceed the maximum file size that the file system allows. You can enter a value from 20 MBs to 51200 MBs. For a FlashBackup policy, Symantec recommends that you use the default, maximum fragment size to ensure optimal deduplication performance.

Table 5-5 Deduplication storage unit properties (*continued*)

Property	Description
Maximum concurrent jobs	<p>The Maximum concurrent jobs setting specifies the maximum number of jobs that NetBackup can send to a disk storage unit at one time. (Default: one job. The job count can range from 0 to 256.) This setting corresponds to the Maximum concurrent write drives setting for a Media Manager storage unit.</p> <p>NetBackup queues jobs until the storage unit is available. If three backup jobs are scheduled and Maximum concurrent jobs is set to two, NetBackup starts the first two jobs and queues the third job. If a job contains multiple copies, each copy applies toward the Maximum concurrent jobs count.</p> <p>Maximum concurrent jobs controls the traffic for backup and duplication jobs but not restore jobs. The count applies to all servers in the storage unit, not per server. If you select multiple media servers in the storage unit and 1 for Maximum concurrent jobs, only one job runs at a time.</p> <p>The number to enter depends on the available disk space and the server's ability to run multiple backup processes.</p> <p>Warning: A Maximum concurrent jobs setting of 0 disables the storage unit.</p>

Deduplication storage unit recommendations

You can use storage unit properties to control how NetBackup performs.

Configure a client-to-server ratio

For a favorable client-to-server ratio, you can use one disk pool and configure multiple storage units to separate your backup traffic. Because all storage units use the same disk pool, you do not have to partition the storage.

For example, assume that you have 100 important clients, 500 regular clients, and four media servers. You can use two media servers to back up your most important clients and two media servers to back up your regular clients.

The following example describes how to configure a favorable client-to-server ratio:

- Configure the media servers for NetBackup deduplication and configure the storage.
- Configure a disk pool.

- Configure a storage unit for your most important clients (such as STU-GOLD). Select the disk pool. Select **Only use the following media servers**. Select two media servers to use for your important backups.
- Create a backup policy for the 100 important clients and select the STU-GOLD storage unit. The media servers that are specified in the storage unit move the client data to the deduplication storage server.
- Configure another storage unit (such as STU-SILVER). Select the same disk pool. Select **Only use the following media servers**. Select the other two media servers.
- Configure a backup policy for the 500 regular clients and select the STU-SILVER storage unit. The media servers that are specified in the storage unit move the client data to the deduplication storage server.

Backup traffic is routed to the wanted data movers by the storage unit settings.

Note: NetBackup uses storage units for media server selection for write activity (backups and duplications) only. For restores, NetBackup chooses among all media servers that can access the disk pool.

Throttle traffic to the media servers

You can use the **Maximum concurrent jobs** settings on disk pool storage units to throttle the traffic to the media servers. Effectively, this setting also directs higher loads to specific media servers when you use multiple storage units for the same disk pool. A higher number of concurrent jobs means that the disk can be busier than if the number is lower.

For example, two storage units use the same set of media servers. One of the storage units (STU-GOLD) has a higher **Maximum concurrent jobs** setting than the other (STU-SILVER). More client backups occur for the storage unit with the higher **Maximum concurrent jobs** setting.

Enabling deduplication encryption

Two procedures exist to enable encryption during deduplication, as follows:

- You can enable encryption on all hosts that deduplicate their own data without configuring them individually.
Use this procedure if you want all of your clients that deduplicate their own data to encrypt that data.
See [“To enable encryption on all hosts”](#) on page 88.
- You can enable encryption on individual hosts.

Use this procedure to enable compression or encryption on the storage server, on a load balancing server, or on a client that deduplicates its own data. See [“To enable encryption on a single host”](#) on page 88.

See [“About deduplication encryption”](#) on page 34.

To enable encryption on all hosts

- 1 On the storage server, open the `contentrouter.cfg` file in a text editor; it resides in the following directory:

```
storage_path/etc/puredisk/contentrouter.cfg
```

- 2 Add `agent_crypt` to the `ServerOptions` line of the file. The following line is an example:

```
ServerOptions=fast,verify_data_read,agent_crypt
```

- 3 If you use load balancing servers, make the same edits to the `contentrouter.cfg` files on those hosts.

To enable encryption on a single host

- 1 Use a text editor to open the `pd.conf` file on the host.

The `pd.conf` file resides in the following directories:

- (UNIX) `/usr/opensv/lib/ost-plugins/`
- (Windows) `install_path\Veritas\NetBackup\bin\ost-plugins`

See [“pd.conf file settings for NetBackup deduplication”](#) on page 119.

- 2 For the line in the file that contains `ENCRYPTION`, remove the pound character (#) in column 1 from that line.
- 3 In that line, replace the 0 (zero) with a 1.

Note: The spaces to the left and right of the equal sign (=) in the file are significant. Ensure that the space characters appear in the file after you edit the file.

- 4 Ensure that the `LOCAL_SETTINGS` parameter is set to 1.

If `LOCAL_SETTINGS` is 0 (zero) and the `ENCRYPTION` setting on the storage server is 0, the client setting does not override the server setting. Consequently, the data is not encrypted on the client host.

- 5 Save and close the file.
- 6 If the host is the storage server or a load balancing server, restart the NetBackup Remote Manager and Monitor Service (`nbrmms`) on the host.

Configuring optimized synthetic backups for deduplication

The following table shows the procedures that are required to configure optimized synthetic backups for a deduplication environment.

See [“About optimized synthetic backups and deduplication”](#) on page 35.

Table 5-6 To configure optimized synthetic backups for deduplication

Step	Task	Procedure
Step 1	Set the OptimizedImage attribute on the storage server.	See “Setting deduplication storage server attributes” on page 152.
Step 2	Set the OptimizedImage attribute on your existing deduplication pools. (Any deduplication pools that you create after you set the storage server attribute inherit the new functionality.)	See “Setting deduplication disk pool attributes” on page 164.
Step 3	Configure a Standard or MS-Windows backup policy. Select the Synthetic backup attribute on the Schedule Attribute tab of the backup policy.	See the administrator's guide for your operating system: <ul style="list-style-type: none"> ■ <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I.</i> ■ <i>NetBackup Administrator's Guide for Windows, Volume I.</i>

See [“Setting deduplication storage server attributes”](#) on page 152.

See [“Creating a policy without using the Policy Configuration Wizard”](#) on page 111.

About configuring optimized duplication and replication bandwidth

Each optimized duplication or Auto Image Replication job is a separate process or stream. The number of duplication or replication jobs that run concurrently determines the number of jobs that contend for bandwidth. You can control how much network bandwidth that optimized duplication and Auto Image Replication jobs consume.

Two different configuration file settings control the bandwidth that is used, as follows:

<code>bandwidthlimit</code>	<p>The <code>bandwidthlimit</code> parameter in the <code>agent.cfg</code> file is the global bandwidth setting. You can use this parameter to limit the bandwidth that all replication jobs use.</p> <p>If <code>bandwidthlimit</code> is greater than zero, all of the jobs share the bandwidth. That is, the bandwidth for each job is the <code>bandwidthlimit</code> divided by the number of jobs.</p> <p>If <code>bandwidthlimit=0</code>, total bandwidth is not limited. However, you can limit the bandwidth that each job uses. See the following <code>OPTDUP_BANDWIDTH</code> description.</p> <p>By default, <code>bandwidthlimit=0</code>.</p>
<code>OPTDUP_BANDWIDTH</code>	<p>The <code>OPTDUP_BANDWIDTH</code> parameter in the <code>pd.conf</code> file specifies the per job bandwidth.</p> <p><code>OPTDUP_BANDWIDTH</code> applies only if the <code>bandwidthlimit</code> parameter in the <code>agent.cfg</code> file is zero.</p> <p>If <code>OPTDUP_BANDWIDTH</code> and <code>bandwidthlimit</code> are both 0, bandwidth per replication job is not limited.</p> <p>By default, <code>OPTDUP_BANDWIDTH = 0</code>.</p>

To specify the bandwidth, edit the configuration files on the deduplication storage server, as follows:

- `bandwidthlimit`.

The `agent.cfg` file resides in the following directory:

- **UNIX:** `storage_path/etc/puredisk`
- **Windows:** `storage_path\etc\puredisk`

- `OPTDUP_BANDWIDTH`.

See [“About the `pd.conf` configuration file for NetBackup deduplication”](#) on page 119.

See [“Editing the `pd.conf` deduplication file”](#) on page 126.

See [“`pd.conf` file settings for NetBackup deduplication ”](#) on page 119.

If you specify bandwidth limits, optimized duplication and replication traffic to any destination is limited.

See [“About MSDP optimized duplication within the same domain”](#) on page 36.

See [“About NetBackup Auto Image Replication”](#) on page 47.

Configuring MSDP optimized duplication copy behavior

You can configure the following optimized duplication behaviors for NetBackup deduplication:

Optimized duplication failover	<p>By default, if an optimized duplication job fails, NetBackup does not run the job again.</p> <p>You can configure NetBackup to use normal duplication if an optimized duplication fails.</p> <p>For example, NetBackup does not support optimized duplication from PureDisk Storage Pool Authority source to a Media Server Deduplication Pool. Both entities support optimized duplication; however, not in this direction. Therefore, to pull or migrate data out of a PureDisk SPA into a Media Server Deduplication Pool, you must change the default NetBackup failover behavior for optimized duplication.</p> <p>See “To configure optimized duplication failover” on page 92.</p>
Number of optimized duplication attempts	<p>By default, NetBackup tries an optimized duplication job three times before it fails the job.</p> <p>You can change the number of times NetBackup retries an optimized duplication job before it fails the jobs.</p> <p>See “To configure the number of duplication attempts” on page 92.</p>
Storage lifecycle policy retry wait period	<p>If a storage lifecycle policy optimized duplication job fails, NetBackup waits two hours and then retries the job. By default, NetBackup tries a job three times before the job fails.</p> <p>You can change the number of hours for the wait period.</p> <p>See “To configure the storage lifecycle policy wait period” on page 92.</p>

Caution: These settings affect all optimized duplication jobs; they are not limited to optimized duplication to a **Media Server Deduplication Pool** or a **PureDisk Deduplication Pool**.

To configure optimized duplication failover

- ◆ On the master server, add the following configuration option:

```
RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE
```

See [“Setting NetBackup configuration options by using bpssetconfig”](#) on page 134.

Alternatively on UNIX systems, add the entry to the `bp.conf` file on the NetBackup master server.

To configure the number of duplication attempts

- ◆ Create a file named `OPT_DUP_BUSY_RETRY_LIMIT` that contains an integer that specifies the number of times to retry the job before NetBackup fails the job.

The file must reside on the master server in the following directory (depending on the operating system):

- UNIX: `/usr/opensv/netbackup/db/config`
- Windows: `install_path\NetBackup\db\config`.

To configure the storage lifecycle policy wait period

- ◆ Change the wait period for retries by adding an `IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS` entry to the NetBackup `LIFECYCLE_PARAMETERS` file. The default for this value is two hours.

For example, the following entry configures NetBackup to wait four hours before NetBackup tries the job again:

```
IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS 4
```

The `LIFECYCLE_PARAMETERS` file resides in the following directories:

- UNIX: `/usr/opensv/netbackup/db/config`
- Windows: `install_path\NetBackup\db\config`.

Configuring a separate network path for MSDP optimized duplication

You can use a separate network for MSDP optimized duplication traffic within the same NetBackup domain.

The following are the requirements:

- A separate, dedicated network interface card in both the source and the destination storage servers.

- The separate network is operational and using the dedicated network interface cards on the source and the destination storage servers.

See “[About MSDP optimized duplication within the same domain](#)” on page 36.

To configure a separate network path for MSDP optimized duplication

- 1 On the source storage server, add the destination storage servers's dedicated network interface to the operating system `hosts` file. If **StorageServer-A** is the source MSDP and **StorageServer-B** is the destination MSDP, the following is an example of the hosts entry in IPv4 notation:

```
192.168.0.0          StorageServer-B.symantecs.org
```

Symantec recommends that you always use the fully qualified domain name when you specify hosts.

- 2 If the source storage server is a UNIX or Linux computer, ensure that it checks the `hosts` file first when resolving host names. To do so, verify that the `/etc/nsswitch.conf` file look-up order is first `files` and then `dns`, as in the following example:

```
hosts:                files dns
```

If the look-up order is different, edit the `/etc/nsswitch.conf` file as necessary.

- 3 On the destination storage server, add the source storage servers's dedicated network interface to the operating system `hosts` file. If **StorageServer-A** is the source MSDP and **StorageServer-B** is the destination MSDP, the following is an example of the hosts entry in IPv4 notation:

```
192.168.0.1          StorageServer-A.symantecs.org
```

Symantec recommends that you always use the fully qualified domain name when specifying hosts.

- 4 If the destination storage server is a UNIX or Linux computer, ensure that it checks the `hosts` file first when resolving host names. To do so, verify that the `/etc/nsswitch.conf` file look-up order is first `files` and then `dns`, as in the following example:

```
hosts:                files dns
```

If the look-up order is different, edit the `/etc/nsswitch.conf` file as necessary.

- 5 From each host, use the `ping` command to verify that the host resolves the name of the host.

```
StorageServer-A.symantecs.org> ping StorageServer-B.symantecs.org
StorageServer-B.symantecs.org> ping StorageServer-A.symantecs.org
```

If the `ping` command returns positive results, the hosts are configured for optimized duplication over the separate network.

Configuring optimized duplication of deduplicated data

You can configure optimized duplication of deduplicated backups. Before you begin, review the requirements.

See [“About MSDP optimized duplication within the same domain”](#) on page 36.

Table 5-7 To configure optimized duplication of deduplicated data

Step	Action	Description
Step 1	Configure the storage servers	<p>One server must be common between the source storage and the destination storage. Which you choose depends on whether you want a push or a pull configuration.</p> <p>See “About the media servers for optimized MSDP duplication within the same domain” on page 38.</p> <p>For a push configuration, configure the common server as a load balancing server for the storage server for your normal backups. For a pull configuration, configure the common server as a load balancing server for the storage server for the copies at your remote site. Alternatively, you can add a server later to either environment. (A server becomes a load balancing server when you select it in the storage unit for the deduplication pool.)</p> <p>See “Optimized MSDP duplication within the same domain requirements” on page 37.</p> <p>See “Configuring a NetBackup deduplication storage server” on page 79.</p>

Table 5-7 To configure optimized duplication of deduplicated data (*continued*)

Step	Action	Description
Step 2	Configure the deduplication pools	<p>If you did not configure the deduplication pools when you configured the storage servers, use the Disk Pool Configuration Wizard to configure them.</p> <p>See “Configuring a deduplication disk pool” on page 81.</p>
Step 3	Configure the storage unit for backups	<p>In the storage unit for your backups, do the following:</p> <ul style="list-style-type: none"> ■ For the Disk type, select PureDisk. ■ For the Disk pool, select one of the following: <ul style="list-style-type: none"> ■ If you back up to integrated NetBackup deduplication, select your Media Server Deduplication Pool. ■ If you back up to a PureDisk environment, select the PureDisk Deduplication Pool. <p>If you use a pull configuration, do not select the common media server in the backup storage unit. If you do, NetBackup uses it to deduplicate backup data. (That is, unless you want to use it for a load balancing server for the source deduplication node.)</p> <p>See “Configuring a deduplication storage unit” on page 83.</p>
Step 4	Configure the storage unit for duplication	<p>Symantec recommends that you configure a storage unit specifically to be the target for the optimized duplication. Configure the storage unit in the deduplication node that performs your normal backups. Do not configure it in the node that contains the copies.</p> <p>In the storage unit that is the destination for your duplicated images, do the following:</p> <ul style="list-style-type: none"> ■ For the Disk type, select PureDisk. ■ For the Disk pool, the destination can be a Media Server Deduplication Pool or a PureDisk Deduplication Pool. <p>Note: If the backup destination is a PureDisk Deduplication Pool, the duplication destination also must be a PureDisk Deduplication Pool.</p> <p>Also select Only use the following media servers. Then, select the media server or media servers that are common to both the source storage server and the destination storage server. If you select more than one, NetBackup assigns the duplication job to the least busy media server.</p> <p>If you select only a media server (or servers) that is not common, the optimized duplication job fails.</p> <p>See “Configuring a deduplication storage unit” on page 83.</p>

Table 5-7 To configure optimized duplication of deduplicated data (*continued*)

Step	Action	Description
Step 5	Configure optimized duplication behaviors	See “Configuring MSDP optimized duplication copy behavior” on page 91. See “About configuring optimized duplication and replication bandwidth” on page 89.
Step 6	Configure a separate network path for optimized duplication traffic	Optionally, you can use a separate network for the optimized duplication traffic. See “Configuring a separate network path for MSDP optimized duplication” on page 92.
Step 7	Configure a storage lifecycle policy for the duplication	<p>Configure a storage lifecycle policy only if you want to use one to duplicate images. The storage lifecycle policy manages both the backup jobs and the duplication jobs. Configure the lifecycle policy in the deduplication environment that performs your normal backups. Do not configure it in the environment that contains the copies.</p> <p>When you configure the storage lifecycle policy, do the following:</p> <ul style="list-style-type: none"> ■ For the Backup destination, select the storage unit that is the target of your backups. That storage unit may use a Media Server Deduplication Pool or a PureDisk Deduplication Pool. These backups are the primary backup copies; they are the source images for the duplication operation. ■ For the Duplication destination, select the storage unit for the destination deduplication pool. That pool may be a Media Server Deduplication Pool or a PureDisk Deduplication Pool. If the backup destination is a PureDisk Deduplication Pool, the duplication destination also must be a PureDisk Deduplication Pool. <p>See “Creating a storage lifecycle policy” on page 105.</p>
Step 8	Configure a backup policy	<p>Configure a policy to back up your clients. Configure the backup policy in the deduplication environment that performs your normal backups. Do not configure it in the environment that contains the copies.</p> <ul style="list-style-type: none"> ■ If you use a storage lifecycle policy to manage the backup job and the duplication job: Select that storage lifecycle policy in the Policy storage field of the Policy Attributes tab. ■ If you do not use a storage lifecycle policy to manage the backup job and the duplication job: Select the storage unit for the Media Server Deduplication Pool that contains your normal backups. These backups are the primary backup copies.

Table 5-7 To configure optimized duplication of deduplicated data (*continued*)

Step	Action	Description
Step 9	Configure NetBackup Vault for the duplication	<p>Configure Vault duplication only if you use NetBackup Vault to duplicate the images. Configure Vault in the deduplication environment that performs your normal backups. Do not configure it in the environment that contains the copies.</p> <p>For Vault, you must configure a Vault profile and a Vault policy.</p> <ul style="list-style-type: none"> ■ Configure a Vault profile. <ul style="list-style-type: none"> ■ On the Vault Profile dialog box Choose Backups tab, choose the backup images in the source Media Server Deduplication Pool. ■ On the Profile dialog box Duplication tab, select the destination storage unit in the Destination Storage unit field. ■ Configure a Vault policy to schedule the duplication jobs. A Vault policy is a NetBackup policy that is configured to run Vault jobs. <p>See the <i>NetBackup Vault Administrator's Guide</i>.</p>
Step 10	Duplicate by using the <code>bpduplicate</code> command	<p>Use the NetBackup <code>bpduplicate</code> command to copy images manually. If you use a storage lifecycle policy or NetBackup Vault for optimized duplication, you do not have to use the <code>bpduplicate</code> command.</p> <p>Duplicate from the source storage to the destination storage. The destination storage may be a Media Server Deduplication Pool or a PureDisk Deduplication Pool.</p> <p>See <i>NetBackup Commands Reference Guide</i>.</p>

About the replication topology for Auto Image Replication

The disk volumes of the devices that support Auto Image Replication have the properties that define the replication relationships between the volumes. The knowledge of the volume properties is considered the replication topology. The following are the replication properties that a volume can have:

- Source** A source volume contains the backups of your clients. The volume is the source for the images that are replicated to a remote NetBackup domain. Each source volume in an originating domain has one or more replication partner target volumes in a target domain.
- Target** A target volume in the remote domain is the replication partner of a source volume in the originating domain.
- None** The volume does not have a replication attribute.

For a **Media Server Deduplication Pool**, NetBackup exposes the storage as a single volume. Therefore, there is always a one-to-one volume relationship for MSDP.

You configure the replication relationships when you add target storage servers in the **Replication** tab of the **Change Storage Server** dialog box.

NetBackup discovers topology changes when you use the **Refresh** option of the **Change Disk Pool** dialog box.

See [“Changing deduplication disk pool properties”](#) on page 165.

NetBackup includes a command that can help you understand your replication topology. Use the command in the following situations:

- After you configure the MSDP replication targets.
- After changes to the volumes that comprise the storage.

See [“Viewing the replication topology for Auto Image Replication”](#) on page 100.

Configuring a target for MSDP replication

Use the following procedure to establish the replication relationship between a **Media Server Deduplication Pool** in an originating domain and a **Media Server Deduplication Pool** in a target domain.

Configuring the target storage server is only one step in the process

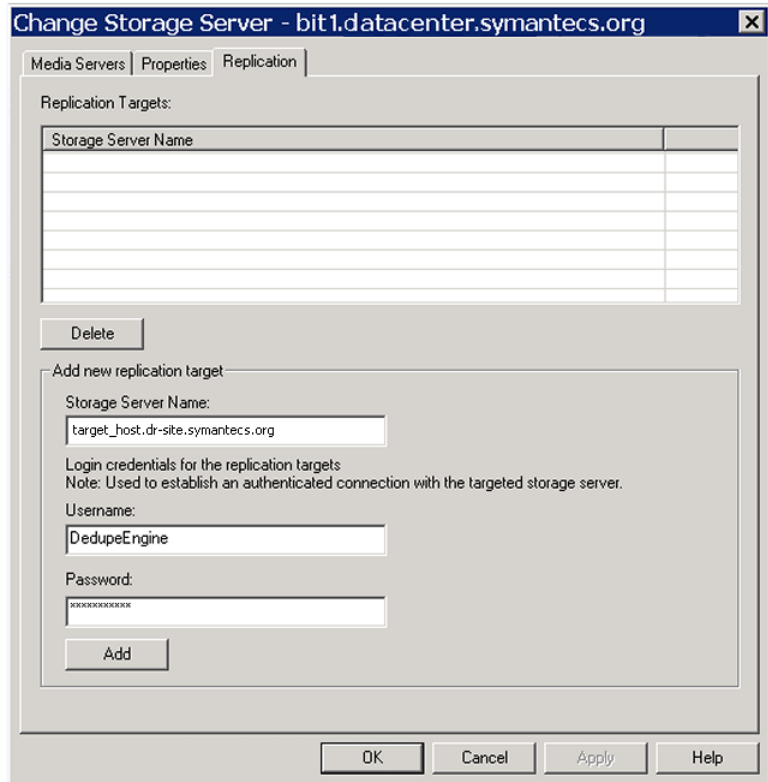
See [“Configuring NetBackup media server deduplication”](#) on page 76.

Caution: Choose the target storage server or servers carefully. A target storage server must not also be a storage server for the source domain.

To configure a Media Server Deduplication Pool as a replication target

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Credentials > Storage Server**.
- 2 Select the MSDP storage server.
- 3 On the **Edit** menu, select **Change**.

- In the **Change Storage Server** dialog box, select the **Replication** tab.



- To add a replication target in a remote domain:
 - Enter the **Storage Server Name**.
 - Enter **Username** and **Password** credentials for the NetBackup Deduplication Engine.
 - Click **Add** to add the storage server to the **Replication Targets** list. After you click **Add**, NetBackup verifies that the target storage server exists. NetBackup also configures the replication properties of the volumes in the source domain and the target domain.

All targets are considered for replication, depending on the rules of the storage lifecycle policies that control the replication.

- 6 After all replication targets are added, click **OK**.
- 7 For the deduplication pool in each domain, open the **Change Disk Pool** dialog box and click **Refresh**.

Configuring a replication target configures the replication properties of the disk volumes in both domains. However, NetBackup only updates the properties of the disk pool when you click **Refresh** in the **Change Disk Pool** dialog box and then click **OK**.

See [“Changing deduplication disk pool properties”](#) on page 165.

Viewing the replication topology for Auto Image Replication

For a replication operation to succeed, a volume that is a source of replication must have at least one replication partner that is the target of replication. NetBackup lets you view the replication topology of the storage.

See [“About the replication topology for Auto Image Replication”](#) on page 97.

To view the replication topology for Auto Image Replication

- ◆ Run the `bpstsinfo` command, specifying the storage server name and the server type. The following is the command syntax:

```
bpstsinfo -lsuinfo -storage_server storage_server_name -stype  
PureDisk
```

The command is located in the following directory:

- UNIX: `/usr/opensv/netbackup/bin/admincmd/`
- Windows: `Install_path\NetBackup\bin\admincmd\`

The following are the options and arguments for the command:

```
-storage_server storage_server_name The name of the storage server.  
  
-stype PureDisk Deduplication storage servers are of type  
PureDisk.
```

Save the output to a file so that you can compare the current topology with the previous topology to determine what has changed.

Example output is available.

See [“Sample volume properties output for MSDP replication”](#) on page 101.

Sample volume properties output for MSDP replication

The following two examples show output from the `bpstsinfo -lsuinfo` command for two NetBackup deduplication storage servers. The first example is the output from the source disk pool in the originating domain. The second example is from the target disk pool in the remote master server domain.

The two examples show the following:

- All of the storage in a deduplication disk pool is exposed as one volume: `PureDiskVolume`.
- The `PureDiskVolume` of the deduplication storage server `bit1.datacenter.symantecs.org` is the source for the replication operation.
- The `PureDiskVolume` of the deduplication storage server `target_host.dr-site.symantecs.org` is the target of the replication operation.

```
> bpstsinfo -lsuinfo -storage_server bit1.datacenter.symantecs.org -stype PureDisk
LSU Info:
    Server Name: PureDisk:bit1.datacenter.symantecs.org
    LSU Name: PureDiskVolume
    Allocation : STS_LSU_AT_STATIC
    Storage: STS_LSU_ST_NONE
    Description: PureDisk storage unit (/bit1.datacenter.symantecs.org#1/2)
    Configuration:
    Media: (STS_LSUF_DISK | STS_LSUF_ACTIVE | STS_LSUF_STORAGE_NOT_FREED |
    STS_LSUF_REP_ENABLED | STS_LSUF_REP_SOURCE)
    Save As : (STS_SA_CLEARF | STS_SA_IMAGE | STS_SA_OPAQUEF)
    Replication Sources: 0 ( )
    Replication Targets: 1 ( PureDisk:target_host.dr-site.symantecs.org:PureDiskVolume )
    Maximum Transfer: 2147483647
    Block Size: 512
    Allocation Size: 0
    Size: 74645270666
    Physical Size: 77304328192
    Bytes Used: 138
    Physical Bytes Used: 2659057664
    Resident Images: 0

> bpstsinfo -lsuinfo -storage_server target_host.dr-site.symantecs.org -stype PureDisk
LSU Info:
    Server Name: PureDisk:target_host.dr-site.symantecs.org
    LSU Name: PureDiskVolume
    Allocation : STS_LSU_AT_STATIC
    Storage: STS_LSU_ST_NONE
```

```
Description: PureDisk storage unit (/target_host.dr-site.symantecs.org#1/2)
Configuration:
Media: (STS_LSU_F_DISK | STS_LSU_F_ACTIVE | STS_LSU_F_STORAGE_NOT_FREED |
STS_LSU_F_REP_ENABLED | STS_LSU_F_REP_TARGET)
Save As : (STS_SA_CLEARF | STS_SA_IMAGE | STS_SA_OPAQUEF)
Replication Sources: 1 ( PureDisk:bit1:PureDiskVolume )
Replication Targets: 0 ( )
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 79808086154
Physical Size: 98944983040
Bytes Used: 138
Physical Bytes Used: 19136897024
Resident Images: 0
```

About the storage lifecycle policies required for Auto Image Replication

To replicate images from the one NetBackup domain to another NetBackup domain requires that two storage lifecycle policies be configured:

- In the first (originating) NetBackup domain:
One SLP that contains at least one **Backup** operation and one **Replication** operation that is configured to replicate to a target NetBackup domain. (The Auto Image Replication SLP.)
- In the second, target NetBackup domain:
One SLP that contains an **Import** operation to import the replication. (The Import SLP.) The Import SLP can be configured to create additional copies in that domain or to cascade the copies to another domain.

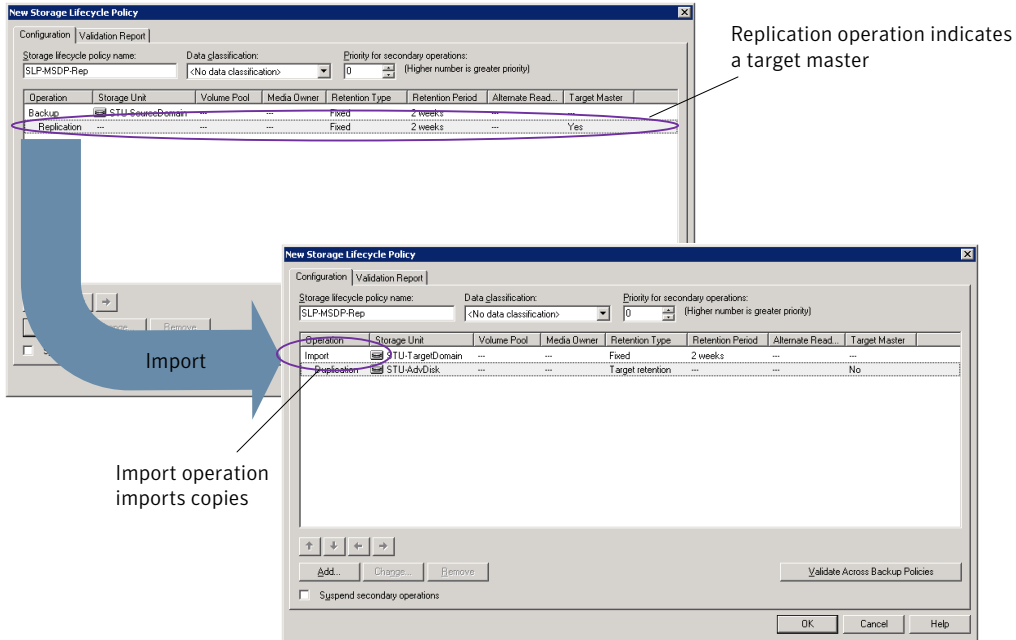
Note: Both SLPs must have identical names.

[Figure 5-1](#) shows how the SLP in the target domain is set up to replicate the images from the originating master server domain.

About the storage lifecycle policies required for Auto Image Replication

Figure 5-1 Storage lifecycle policy pair required for Auto Image Replication

SLP on master server in the source domain



SLP that imports the copies to the target domain

Table 5-8 describes the requirements for each SLP in the pair.

Table 5-8 SLP requirements for Auto Image Replication

Domain	Storage lifecycle policy requirements
Domain 1 (Originating domain)	<p>The Auto Image Replication SLP must meet the following criteria:</p> <ul style="list-style-type: none"> ■ The SLP must have the same name as the Import SLP in Domain 2. ■ The SLP must be of the same data classification as the Import SLP in Domain 2. ■ The Backup operation must be to a Media Server Deduplication Pool (MSDP). Indicate the exact storage unit from the drop-down list. Do not select Any Available. Note: The target domain must contain the same type of storage to import the image. ■ At least one operation must be a Replication operation with the Target master option selected. See Figure 5-2 on page 104. <p>Multiple Replication operations can be configured in an Auto Image Replication SLP. The master server in Domain 1 does not know which target media server will be selected. If multiple SLPs in target domains meet the criteria, NetBackup imports copies in all qualifying domains.</p>

Table 5-8 SLP requirements for Auto Image Replication (*continued*)

Domain	Storage lifecycle policy requirements
Domain 2 (Target domain)	<p>The Import SLP must meet the following criteria:</p> <ul style="list-style-type: none"> ■ The SLP must have the same name as the SLP in Domain 1 described above. The matching name indicates to the SLP which images to process. ■ The SLP must be of the same data classification as the SLP in Domain 1 described above. Matching the data classification keeps a consistent meaning to the classification and facilitates global reporting by data classification. ■ The first operation in the SLP must be an Import operation. Indicate the exact storage unit from the drop-down list. Do not select Any Available. See Figure 5-3 on page 105. ■ The SLP must contain at least one Replication operation that has the Target retention specified.

The following topic describes useful reporting information about Auto Image Replication jobs and import jobs.

See “[Reporting on Auto Image Replication jobs](#) ” on page 147.

Figure 5-2 Replication operation with Target master option selected in Domain 1 storage lifecycle policy

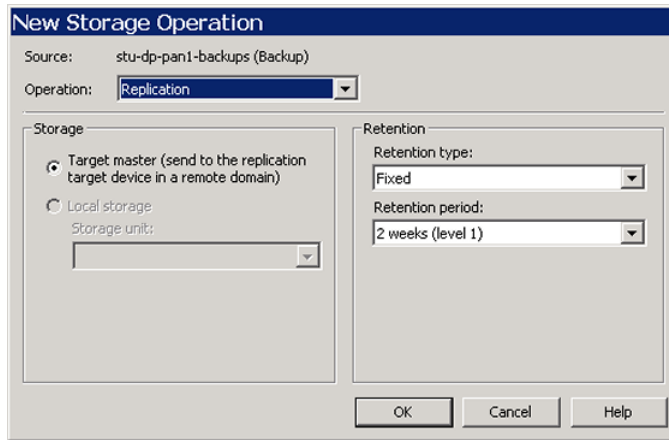



Figure 5-3 Import operation in Domain 2 storage lifecycle policy



See “[Creating a storage lifecycle policy](#)” on page 105.

Customizing how nbstserv runs duplication and import jobs

The NetBackup Storage Lifecycle Manager (*nbstserv*) runs replication, duplication, and import jobs. Both the Duplication Manager service and the Import Manager service run within *nbstserv*.

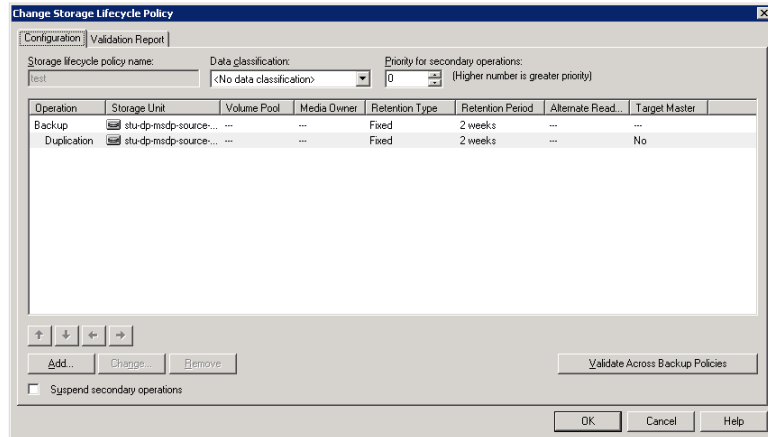
The NetBackup administrator can customize how *nbstserv* runs jobs by adding parameters to the `LIFECYCLE_PARAMETERS` file.

Creating a storage lifecycle policy

A storage lifecycle policy can be selected as the **Policy storage** within a backup policy.

To create a storage lifecycle policy

- 1 In the **NetBackup Administration Console**, select **NetBackup Management > Storage > Storage Lifecycle Policies**.
- 2 Click **Actions > New > Storage Lifecycle Policy (UNIX)** or **Actions > New > New Storage Lifecycle Policy (Windows)**.



- 3 In the **New Storage Lifecycle Policy** dialog box, enter a **Storage lifecycle policy name**.
- 4 Select a **Data classification**. (Optional.)
- 5 Select the **Priority for secondary operations**. This number represents the priority that jobs from secondary operations have in relationship to all other jobs.
See [“Storage Lifecycle Policy dialog box settings”](#) on page 106.
- 6 Click **Add** to add operations to the SLP. The operations act as instructions for the data.
See [“Adding a storage operation to a storage lifecycle policy”](#) on page 108.
- 7 Click **OK** to create the storage lifecycle policy.

Storage Lifecycle Policy dialog box settings

A storage lifecycle policy consists of one or more operations.

The **New Storage Lifecycle** dialog box and the **Change Storage Lifecycle Policy** dialog box contains the following settings.

Figure 5-4 Configuration tab of the Storage Lifecycle Policy dialog box

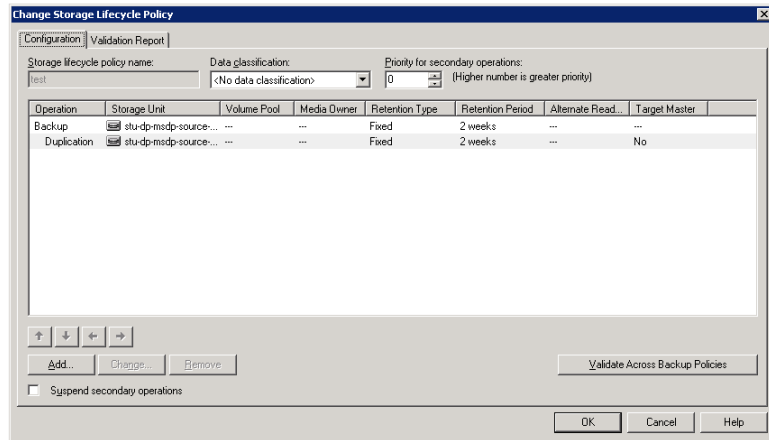


Table 5-9 Configuration tab of the Storage Lifecycle Policy dialog box

Setting	Description
Storage lifecycle policy name	The Storage lifecycle policy name describes the SLP. The name cannot be modified after the SLP is created.
Data classification	<p>The Data classification defines the level of data that the SLP is allowed to process. The Data classification drop-down menu contains all of the defined classifications. The Data classification is an optional setting.</p> <p>One data classification can be assigned to each SLP and applies to all operations in the SLP. An SLP is not required to have a data classification.</p> <p>If a data classification is selected, the SLP stores only those images from the policies that are set up for that data classification. If no data classification is indicated, the SLP accepts images of any classification or no classification.</p> <p>The Data classification setting allows the NetBackup administrator to classify data based on relative importance. A classification represents a set of backup requirements. When data must meet different backup requirements, consider assigning different classifications.</p> <p>For example, email backup data can be assigned to the silver data classification and financial data backup may be assigned to the platinum classification.</p> <p>A backup policy associates backup data with a data classification. Policy data can be stored only in an SLP with the same data classification.</p> <p>Once data is backed up in an SLP, the data is managed according to the SLP configuration. The SLP defines what happens to the data from the initial backup until the last copy of the image has expired.</p>

Table 5-9 Configuration tab of the Storage Lifecycle Policy dialog box
(continued)

Setting	Description
Priority for secondary operations	<p>The Priority for secondary operations setting is the priority that secondary jobs (for example, duplication jobs), have in relationship to all other jobs. Range: 0 (default) to 99999 (highest priority).</p> <p>For example, the Priority for secondary operations for a policy with a gold data classification may be set higher than for a policy with a silver data classification.</p> <p>The priority of the backup job is set in the backup policy on the Attributes tab.</p>
Operations	<p>The Operations list contains all of the operations in the SLP. Multiple operations imply that multiple copies are created.</p> <p>The list also contains the columns that display information about each operation. Note that not all columns display by default.</p> <p>For column descriptions, see the following topic:</p>
Suspend secondary operations	<p>Enable Suspend secondary operations to stop the operations in the SLP.</p> <p>A selected SLP can also be suspended from the Actions menu and then activated again (Activate).</p>
Validate Across Backup Policies button	<p>Use this button to see how changes to this SLP can affect the policies that are associated with this SLP. The button generates a report that displays on the Validation Report tab.</p> <p>This button performs the same validation as the <code>-conflict</code> option performs when used with the <code>nbstl</code> command.</p>
Arrows	<p>Use the arrows to indicate the indentation (or hierarchy) of the source for each copy. One copy can be the source for many other copies.</p> <p>Many operations can be hierarchical or non-hierarchical:</p>

Adding a storage operation to a storage lifecycle policy

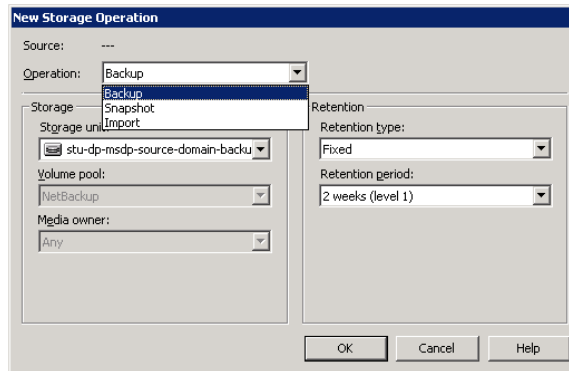
Use the following procedure to add a storage operation to a storage lifecycle policy:

To add a storage operation to a lifecycle policy

- 1 In the **NetBackup Administration Console**, select **NetBackup Management > Storage > Storage Lifecycle Policies**.
- 2 Click **Actions > New > New Storage Lifecycle Policy (Windows)** or **Actions > New > Storage Lifecycle Policy (UNIX)**.

- 3 Click **Add** to add operations to the SLP. The operations are the instructions for the SLP to follow and apply to the data that is eventually specified in the backup policy.

To create a hierarchical SLP, select an operation to become the source of the next operation, then click **Add**.



- 4 In the **New Storage Operation** dialog box, select an **Operation** type. The name of the operation reflects its purpose in the SLP:

- **Backup**
- **Backup From Snapshot**
- **Duplication**
 See “[About NetBackup Auto Image Replication](#)” on page 47.
- **Import**
- **Index From Snapshot**
- **Replication**
- **Snapshot**

- 5 Indicate where the operation is to write the image. Depending on the operation, selections may include storage units or storage unit groups.

No BasicDisk, SnapVault, or disk staging storage units can be used as storage unit selections in an SLP.

Note: In NetBackup 7.5, the **Any_Available** selection is not available for new SLPs. In an upgrade situation, existing SLPs that use **Any_Available** continue to work as they did before NetBackup 7.5. However, if the NetBackup administrator edits an existing SLP, a specific storage unit or storage unit group must be selected before the SLP can be saved successfully.

- 6 If the storage unit is a tape device or virtual tape library (VTL), indicate the **Volume pool** where the backups (or copies) are to be written.
- 7 Indicate the **Media owner** if the storage unit is a Media Manager type and server groups are configured.

By specifying a **Media owner**, you allow only those media servers to write to the media on which backup images for this policy are written.
- 8 Select the retention type for the operation:
 - **Capacity managed**
 - **Expire after copy**
If a policy is configured to back up to a lifecycle, the retention that is indicated in the lifecycle is the value that is used. The **Retention** attribute in the schedule is not used.
 - **Fixed**
 - **Maximum snapshot limit**
 - **Mirror**
 - **Target retention**
- 9 Indicate an **Alternate read server** that is allowed to read a backup image originally written by a different media server.
- 10 Click **OK** to create the storage operation.

About backup policy configuration

When you configure a backup policy, for the **Policy storage** select a storage unit that uses a deduplication pool.

For a storage lifecycle policy, for the **Storage unit** select a storage unit that uses a deduplication pool.

For VMware backups, select the **Enable file recovery from VM backup** option when you configure a VMware backup policy. The **Enable file recovery from VM backup** option provides the best deduplication rates.

NetBackup deduplicates the client data that it sends to a deduplication storage unit.

Creating a policy using the Policy Configuration Wizard

The easiest method to set up a backup policy is to use the **Policy Configuration Wizard**. This wizard guides you through the setup process by automatically choosing the best values for most configurations.

Not all policy configuration options are presented through the wizard. For example, calendar-based scheduling and the **Data Classification** setting. After the policy is created, modify the policy in the **Policies** utility to configure the options that are not part of the wizard.

Use the following procedure to create a policy using the Policy Configuration Wizard.

To create a policy with the Policy Configuration Wizard

- 1 In the **NetBackup Administration Console**, in the left pane, click **NetBackup Management**.
- 2 In the right pane, click **Create a Policy** to begin the **Policy Configuration Wizard**.
- 3 Select **File systems, databases, or applications**.
- 4 Click **Next** to start the wizard and follow the prompts.

Click **Help** on any wizard panel for assistance while running the wizard.

Creating a policy without using the Policy Configuration Wizard

Use the following procedure to create a policy without using the Policy Configuration Wizard.

To create a policy without the Policy Configuration Wizard

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Type a unique name for the new policy in the **Add a New Policy** dialog box. See “[NetBackup naming conventions](#)” on page 22.
- 3 If necessary, clear the **Use Policy Configuration Wizard** checkbox.
- 4 Click **OK**.
- 5 Configure the attributes, the schedules, the clients, and the backup selections for the new policy.

Enabling client-side deduplication

To enable client deduplication, set an attribute in the NetBackup master server **Client Attributes** host properties. If the client is in a backup policy in which the storage destination is a **Media Server Deduplication Pool**, the client deduplicates its own data.

To specify the clients that deduplicate backups

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the details pane, select the master server.
- 3 On the **Actions** menu, select **Properties**.
- 4 On the **Host Properties General** tab, add the clients that use client direct to the **Clients** list.
- 5 Select one of the following **Deduplication Location** options:
 - **Always use the media server** disables client deduplication. By default, all clients are configured with the **Always use the media server** option.
 - **Prefer to use client-side deduplication** uses client deduplication if the deduplication plug-in is active on the client. If it is not active, a normal backup occurs; client deduplication does not occur.
 - **Always use client-side deduplication** uses client deduplication. If the deduplication backup job fails, NetBackup retries the job.

You can override the **Prefer to use client-side deduplication** or **Always use client-side deduplication** host property in the backup policies.

See **Disable client-side deduplication** in the *NetBackup Administrator's Guide for UNIX and Linux, Volume I*.

See **Disable client-side deduplication** in the *NetBackup Administrator's Guide for Windows, Volume I*.

Resilient Network properties

The **Resilient Network** properties appear for the master server, for media servers, and for clients. For media servers and clients, the **Resilient Network** properties are read only. When a job runs, the master server updates the media server and the client with the current properties.

The **Resilient Network** properties let you configure NetBackup to use resilient network connections. A resilient connection allows backup and restore traffic between a client and NetBackup media servers to function effectively in high-latency, low-bandwidth networks such as WANs. The use case that benefits the most from a resilient connection is a client in a remote office that backs up its own data (client-side deduplication). The data travels across a wide area network (WAN) to media servers in a central datacenter.

NetBackup monitors the socket connections between the remote client and the NetBackup media server. If possible, NetBackup re-establishes dropped connections and resynchronizes the data stream. NetBackup also overcomes latency issues to maintain an unbroken data stream. A resilient connection can survive network interruptions of up to 80 seconds. A resilient connection may survive interruptions longer than 80 seconds.

The NetBackup Remote Network Transport Service manages the connection between the computers. The Remote Network Transport Service runs on the master server, the client, and the media server that processes the backup or restore job. If the connection is interrupted or fails, the services attempt to re-establish a connection and synchronize the data. More information about the Remote Network Transport Service is available.

Resilient connections apply between clients and NetBackup media servers, which includes master servers when they function as media servers. Resilient connections do not apply to master servers or media servers if they function as clients and back up data to a media server.

Resilient connections can apply to all of the clients or to a subset of clients.

Note: If a client is in a different subdomain than the server, add the fully qualified domain name of the server to the client's hosts file. For example, india.symantecs.org is a different subdomain than china.symantecs.org.

When a backup or restore job for a client starts, NetBackup searches the **Resilient Network** list from top to bottom looking for the client. If NetBackup finds the

client, NetBackup updates the resilient network setting of the client and the media server that runs the job. NetBackup then uses a resilient connection.

Figure 5-5 Master server Resilient Network host properties

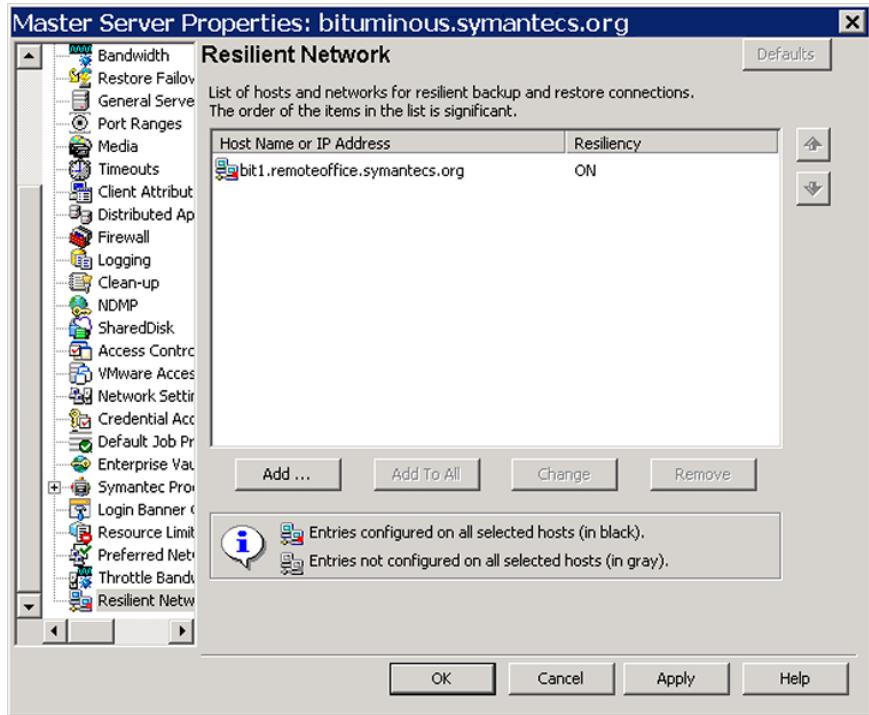


Table 5-10 describes the **Resilient Network** properties.

Table 5-10 Resilient Network dialog box properties

Property	Description
Host Name or IP Address	<p>The Host Name or IP Address of the host. The address can also be a range of IP addresses so you can configure more than one client at once. You can mix IPv4 addresses and ranges with IPv6 addresses and subnets.</p> <p>If you specify the host by name, Symantec recommends that you use the fully qualified domain name.</p> <p>Use the arrow buttons on the right side of the pane to move up or move down an item in the list of resilient networks.</p>
Resiliency	Resiliency is either ON or OFF .

Note: The order is significant for the items in the list of resilient networks. If a client is in the list more than once, the first match determines its resilient connection status. For example, suppose you add a client and specify the client IP address and specify **On** for **Resiliency**. Suppose also that you add a range of IP addresses as **Off**, and the client IP address is within that range. If the client IP address appears before the address range, the client connection is resilient. Conversely, if the IP range appears first, the client connection is not resilient.

The resilient status of each client also appears as follows:

- In the **NetBackup Administration Console**, select **NetBackup Management > Policies** in the left pane and then select a policy. In the right pane, a **Resiliency** column shows the status for each client in the policy.
- In the **NetBackup Administration Console**, select **NetBackup Management > Host Properties > Clients** in the left pane. In the right pane, a **Resiliency** column shows the status for each client.

Other NetBackup properties control the order in which NetBackup uses network addresses.

The NetBackup resilient connections use the SOCKS protocol version 5.

Resilient connection traffic is not encrypted. Symantec recommends that you encrypt your backups. For deduplication backups, use the deduplication-based encryption. For other backups, use policy-based encryption.

Resilient connections apply to backup connections. Therefore, no additional network ports or firewall ports must be opened.

Resilient connection resource usage

Resilient connections consume more resources than regular connections, as follows:

- More socket connections are required per data stream. Three socket connections are required to accommodate the Remote Network Transport Service that runs on both the media server and the client. Only one socket connection is required for a non-resilient connection.
- More sockets are open on media servers and clients. Three open sockets are required rather than one for a non-resilient connection. The increased number of open sockets may cause issues on busy media servers.
- More processes run on media servers and clients. Usually, only one more process per host runs even if multiple connections exist.

- The processing that is required to maintain a resilient connection may reduce performance slightly.

Specifying resilient connections

Use the following procedure to specify resilient connections for NetBackup clients.

See “[Resilient Network properties](#)” on page 113.

Alternatively, you can use the `resilient_clients goodies` script to specify resilient connections for clients:

To specify resilient connections

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers** in the left pane.
- 2 In the right pane, select the host or hosts on which to specify properties.
- 3 Click **Actions > Properties**.
- 4 In the properties dialog box left pane, select **Resilient Network**.
- 5 In the **Resilient Network** dialog box, use the following buttons to manage resiliency:

Add Opens a dialog box in which you can add a host or an address range.
If you specify the host by name, Symantec recommends that you use the fully qualified domain name.

Add To All If you select multiple hosts in the **NetBackup Administration Console**, the entries in the **Resilient Network** list may appear in different colors, as follows:

- The entries that appear in black type are configured on all of the hosts.
- The entries that appear in gray type are configured on some of the hosts only.

For the entries that are configured on some of the hosts only, you can add them to all of the hosts. To do so, select them and click **Add To All**.

Change Opens a dialog box in which you can change the resiliency settings of the select items.

Remove Remove the select host or address range. A confirmation dialog box does not appear.



Move the selected item or items up or down.
The order of the items in the list is significant.
See [“Resilient Network properties”](#) on page 113.

Seeding the fingerprint cache for remote client-side deduplication

Normally, the first backup of a client is more time consuming than successive backups. In a remote client scenario, the communication lag across a WAN adds to the backup time.

See [“Client-side deduplication backup process”](#) on page 220.

You can improve the performance of the first backup of a remote deduplication client. To do so, seed the client's fingerprint cache with the fingerprints from an existing similar client.

To seed the fingerprint cache for remote client-side deduplication

- ◆ Before the first backup of the remote client, edit the `pd.conf` file on the remote client to add the name of the existing similar client. Also add the name of the backup policy for that client and the last date on which to use that client's fingerprint cache.

Specify the setting in the following format:

clienthostmachine,backuppolicy,date

The *date* is the last date in mm/dd/yyyy format to use the fingerprint cache from the client you specify.

See [“Editing the pd.conf deduplication file”](#) on page 126.

See [“pd.conf file settings for NetBackup deduplication ”](#) on page 119.

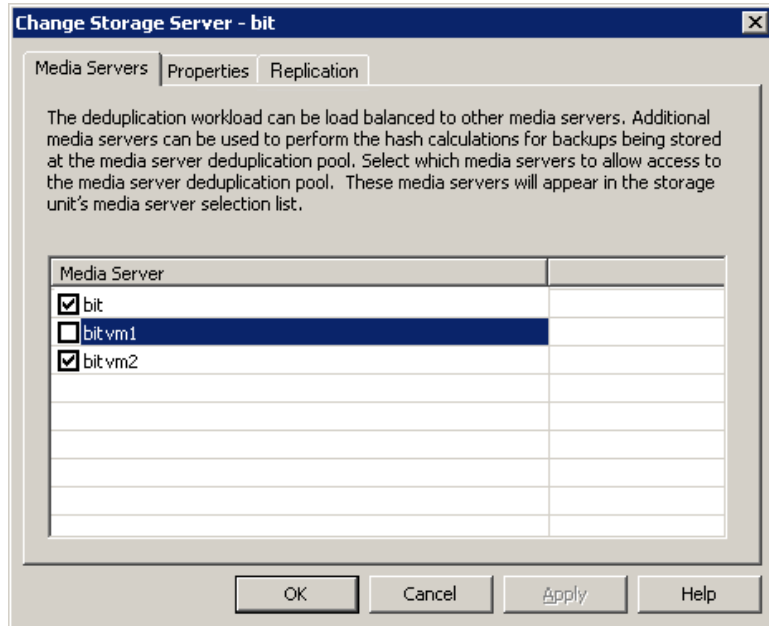
Adding a deduplication load balancing server

You can add a load balancing server to an existing media server deduplication node.

See [“About NetBackup deduplication servers”](#) on page 25.

To add a load balancing server

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Credentials > Storage Server**
- 2 Select the deduplication storage server.
- 3 On the **Edit**, select **Change**.



- 4 In the **Change Storage Server** dialog box, select the **Media Servers** tab
- 5 Select the media server or servers that you want to use as a load balancing server. It must be a supported host.
The media servers that are checked are configured as load balancing servers.
- 6 Click **OK**.
- 7 For all storage units in which **Only use the following media servers** is configured, ensure that the new load balancing server is selected.

About the `pd.conf` configuration file for NetBackup deduplication

On each NetBackup host that deduplicates data, a `pd.conf` file contains the various configuration settings that control the operation of deduplication for the host. By default, the `pd.conf` file settings on the deduplication storage server apply to all load balancing servers and all clients that deduplicate their own data.

You can edit the file to configure advanced settings for that host. If a configuration setting does not exist in a `pd.conf` file, you can add it. If you change the `pd.conf` file on a host, it changes the settings for that host only. If you want the same settings for all of the hosts that deduplicate data, you must change the `pd.conf` file on all of the hosts.

The `pd.conf` file settings may change between releases. During upgrades, NetBackup adds only required settings to existing `pd.conf` files.

The `pd.conf` file resides in the following directories:

- (UNIX) `/usr/opensv/lib/ost-plugins/`
- (Windows) `install_path\Veritas\NetBackup\bin\ost-plugins`

`pd.conf` file settings for NetBackup deduplication

The following table describes the deduplication settings that you can configure.

The parameters in this table are in alphabetical order; the parameters in a `pd.conf` file may not be in alphabetical order.

The settings in your release may differ from the settings that described in this topic.

You can edit the file to configure advanced settings for a host. If a setting does not exist in a `pd.conf` file, you can add it. During upgrades, NetBackup adds only required settings to existing `pd.conf` files.

Table 5-11 `pd.conf` file parameters

Setting	Description
<code>BACKUPRESTORERANGE</code>	<p>On a client, specifies the IP address or range of addresses that the local network interface card (NIC) should use for backups and restores.</p> <p>Specify the value in one of two ways, as follows:</p> <ul style="list-style-type: none"> ■ Classless Inter-Domain Routing (CIDR) format. For example, the following notation specifies 192.168.10.0 and 192.168.10.1 for traffic: <code>BACKUPRESTORERANGE = 192.168.10.1/31</code> ■ Comma-separated list of IP addresses. For example, the following notation specifies 192.168.10.1 and 192.168.10.2 for traffic: <code>BACKUPRESTORERANGE = 192.168.10.1, 192.168.10.2</code> <p>Default value: <code>BACKUPRESTORERANGE=</code> (no default value)</p> <p>Possible values: Classless Inter-Domain Routing format notation or comma-separated list of IP addresses</p>
<code>BANDWIDTH_LIMIT</code>	<p>Determines the maximum bandwidth that is allowed when backing up or restoring data between the deduplication host and the deduplication pool. The value is specified in KBytes/second. The default is no limit.</p> <p>Default value: <code>BANDWIDTH_LIMIT = 0</code></p> <p>Possible values: 0 (no limit) to the practical system limit, in KBs/sec</p>
<code>CLIENT_POLICY_DATE</code>	<p>Note: This setting is valid only in a <code>pd.conf</code> file on a client that deduplicates its own data.</p> <p>Specifies the client, the backup policy, and the date for which to perform client-side rebasing. Specify the client that hosts the <code>pd.conf</code> file.</p> <p>If you specify a value, NetBackup does not check for the existence of segments on the date specified. In the deduplication pool, NetBackup stores these segments next to each other within containers. Queue processing removes the file segments that were stored previously, which maintains unique segments within the deduplication pool.</p> <p>Specify the setting in a <i>clienthostmachine,backuppolicy,date</i> (mm/dd/yyyy) format. The <i>date</i> is the date on which to not check the existence of a segment on the deduplication storage.</p> <p>NetBackup automatically adds the <code>CLIENT_POLICY_DATE</code> parameter to Linux and UNIX NetBackup client-side deduplication clients. If rebasing is required on Windows clients, edit the <code>pd.conf</code> file and add the parameter manually.</p> <p>Default value: <code>CLIENT_POLICY_DATE =</code> (no default value)</p> <p>See “About deduplication storage rebasing” on page 180.</p>

Table 5-11 pd.conf file parameters (*continued*)

Setting	Description
COMPRESSION	<p>Specifies whether to compress the data. By default, files are not compressed.</p> <p>Default value: <code>COMPRESSION = 1</code></p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About deduplication compression” on page 34.</p>
CR_STATS_TIMER	<p>Specifies a time interval in seconds for retrieving statistics from the storage server host. The default value of 0 disables caching and retrieves statistics on demand. Longer intervals may help with larger concurrency.</p> <p>Consider the following information before you change this setting:</p> <ul style="list-style-type: none"> ■ If disabled (set to 0), a request for the latest storage capacity information occurs whenever NetBackup requests it. ■ If you specify a value, a request occurs only after the specified number of seconds since the last request. Otherwise, a cached value from the previous request is used. ■ Enabling this setting may reduce the queries to the storage server. The drawback is the capacity information reported by NetBackup becomes stale. Therefore, if storage capacity is close to full, Symantec recommends that you do not enable this option. ■ On high load systems, the load may delay the capacity information reporting. If so, NetBackup may mark the storage unit as down. <p>Default value: <code>CR_STATS_TIMER = 0</code></p> <p>Possible values: 0 or greater, in seconds</p>
DEBUGLOG	<p>Specifies the file to which NetBackup writes the deduplication log information. NetBackup prepends a date stamp to each day's log file.</p> <p>On Windows, a partition identifier and slash must precede the file name. On UNIX, a slash must precede the file name.</p> <p>Default value:</p> <ul style="list-style-type: none"> ■ UNIX: <code>DEBUGLOG = /var/log/puredisk/pdplugin.log</code> ■ Windows: <code>DEBUGLOG = C:\pdplugin.log C:\pdplugin.log</code> <p>Possible values: Any path</p>

Table 5-11 pd.conf file parameters (*continued*)

Setting	Description
DONT_SEGMENT_TYPES	<p>A comma-separated list of file name extensions of files not to be deduplicated. Files in the backup stream that have the specified extensions are given a single segment if smaller than 16 MB. Larger files are deduplicated using the maximum 16-MB segment size.</p> <p>Example: DONT_SEGMENT_TYPES = mp3,avi.</p> <p>This setting prevents NetBackup from analyzing and managing segments within the file types that do not deduplicate globally.</p> <p>Default value: DONT_SEGMENT_TYPES = (no default value)</p> <p>Possible values: comma-separated file extensions</p>
ENCRYPTION	<p>Specifies whether to encrypt the data. By default, files are not encrypted..</p> <p>If you set this parameter to 1 on all hosts, the data is encrypted during transfer and on the storage.</p> <p>Default value: ENCRYPTION = 0</p> <p>Possible values: 0 (no encryption) or 1 (encryption)</p> <p>See “About deduplication encryption” on page 34.</p>
FIBRECHANNEL	<p>Enable Fibre Channel for backup and restore traffic to and from a NetBackup series appliance.</p> <p>See “About Fibre Channel to a NetBackup 5020 appliance” on page 228.</p> <p>Default value: FIBRECHANNEL = 0</p> <p>Possible values: 0 (off) or 1 (on)</p>
FP_CACHE_LOCAL	<p>Specifies whether or not to use the fingerprint cache for the backup jobs that are deduplicated on the storage server. This parameter does not apply to load balancing servers or to clients that deduplicate their own data.</p> <p>When the deduplication job is on the same host as the NetBackup Deduplication Engine, disabling the fingerprint cache improves performance.</p> <p>Default value: FP_CACHE_LOCAL = 0</p> <p>Possible values: 0 (off) or 1 (on)</p>
FP_CACHE_MAX_MBSIZE	<p>Specifies the amount of memory in MBs to use for the fingerprint cache.</p> <p>Default value: FP_CACHE_MAX_MBSIZE = 20</p> <p>Possible values: 0 to the computer limit</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>

Table 5-11 pd.conf file parameters (*continued*)

Setting	Description
FP_CACHE_MAX_COUNT	<p>Specifies the maximum number of images to load in the fingerprint cache.</p> <p>Default value: FP_CACHE_MAX_COUNT = 1024</p> <p>Possible values: 0 to 4096</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>
FP_CACHE_INCREMENTAL	<p>Specifies whether to use fingerprint caching for incremental backups.</p> <p>Because incremental backups only back up what has changed since the last backup, cache loading has little affect on backup performance for incremental backups.</p> <p>Default value: FP_CACHE_INCREMENTAL = 0</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>
FP_CACHE_CLIENT_POLICY	<p>Note: Symantec recommends that you use this setting on the individual clients that back up their own data (client-side deduplication). If you use it on a storage server or load balancing server, it affects all backup jobs.</p> <p>Specifies the client, backup policy, and date from which to obtain the fingerprint cache for the first backup of a client.</p> <p>By default, the fingerprints from the previous backup are loaded. This parameter lets you load the fingerprint cache from another, similar backup. It can reduce the amount of time that is required for the first backup of a client. This parameter especially useful for remote office backups to a central datacenter in which data travels long distances over a WAN.</p> <p>Specify the setting in the following format:</p> <p><i>clienthostmachine,backuppolicy,date</i></p> <p>The <i>date</i> is the last date in mm/dd/yyyy format to use the fingerprint cache from the client you specify.</p> <p>Default value: FP_CACHE_CLIENT_POLICY = (no default value)</p> <p>See “Seeding the fingerprint cache for remote client-side deduplication” on page 117.</p>

Table 5-11 `pd.conf` file parameters (continued)

Setting	Description
LOCAL_SETTINGS	<p>Specifies whether to allow the <code>pd.conf</code> settings of the deduplication storage server to override the settings in the local <code>pd.conf</code> file.</p> <p>Set this value to 1 if you use local <code>SEGKSIZE</code>, <code>MINFILE_KSIZE</code>, <code>MATCH_PDRO</code>, and <code>DONT_SEGMENT_TYPES</code> settings.</p> <p>Default value: <code>LOCAL_SETTINGS = 0</code></p> <p>Possible values: 0 (allow override) or 1 (always use local settings)</p>
LOGLEVEL	<p>Specifies the amount of information that is written to the log file. The range is from 0 to 10, with 10 being the most logging.</p> <p>Default value: <code>LOGLEVEL = 0</code></p> <p>Possible values: An integer, 0 to 10 inclusive</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>
MAX_IMG_MBSIZE	<p>The maximum backup image fragment size in megabytes.</p> <p>Default value: <code>MAX_IMG_MBSIZE = 51200</code></p> <p>Possible values: 0 to 51,200, in MBs</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>
MAX_LOG_MBSIZE	<p>The maximum size of the log file in megabytes. NetBackup creates a new log file when the log file reaches this limit. NetBackup prepends the date and the ordinal number beginning with 0 to each log file, such as <code>120131_0_pdplugin.log</code>, <code>120131_1_pdplugin.log</code>, and so on.</p> <p>Default value: <code>MAX_LOG_MBSIZE = 500</code></p> <p>Possible values: 0 to 50,000, in MBs</p>
META_SEGKSIZE	<p>The segment size for metadata streams</p> <p>Default value: <code>META_SEGKSIZE = 16384</code></p> <p>Possible values: 32-16384, multiples of 32</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>

Table 5-11 pd.conf file parameters (*continued*)

Setting	Description
OPTDUP_BANDWIDTH	<p>Determines the bandwidth that is allowed for each optimized duplication and Auto Image Replication stream on a deduplication server. OPTDUP_BANDWIDTH does not apply to clients. The value is specified in KBytes/second.</p> <p>Default value: OPTDUP_BANDWIDTH = 0</p> <p>Possible values: 0 (no limit) to the practical system limit, in KBs/sec</p> <p>A global bandwidth parameter affects whether or not OPTDUP_BANDWIDTH applies.</p> <p>See “About configuring optimized duplication and replication bandwidth” on page 89.</p>
OPTDUP_COMPRESSION	<p>Specifies whether to compress optimized duplication data. By default, files are compressed. To disable compression, change the value to 0.</p> <p>Default value: OPTDUP_COMPRESSION = 1</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About deduplication compression” on page 34.</p>
OPTDUP_ENCRYPTION	<p>Specifies whether to encrypt the optimized duplication data. By default, files are not encrypted. If you want encryption, change the value to 1.</p> <p>If you set this parameter to 1 on all hosts, the data is encrypted during transfer and on the storage.</p> <p>Default value: OPTDUP_ENCRYPTION = 0</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About deduplication encryption” on page 34.</p>
OPTDUP_TIMEOUT	<p>Specifies the number of minutes before the optimized duplication times out.</p> <p>Default value: OPTDUP_TIMEOUT = 720</p> <p>Possible values: The value, expressed in minutes</p>
PREFERRED_EXT_SEGKSIZE	<p>Specifies the file extensions and the preferred segment sizes in KB for specific file types. File extensions are case sensitive. The following describe the default values: <code>edb</code> are Microsoft Exchange files; <code>mdf</code> are Microsoft SQL master database files, <code>ndf</code> are Microsoft SQL secondary data files, and <code>segsize64k</code> are Microsoft SQL streams.</p> <p>Default value: PREFERRED_EXT_SEGKSIZE = <code>edb:32,mdf:64,ndf:64,segsize64k:64</code></p> <p>Possible values: <i>file_extension:segment_size_in_KBs</i> pairs, separate by commas.</p> <p>See also SEGKSIZE.</p>

Table 5-11 pd.conf file parameters (*continued*)

Setting	Description
PREFETCH_SIZE	<p>The size in bytes to use for the data buffer for restore operations.</p> <p>Default value: PREFETCH_SIZE = 33554432</p> <p>Possible values: 0 to the computer's memory limit</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>
RESTORE_DECRYPT_LOCAL	<p>Specifies on which host to decrypt and decompress the data during restore operations.</p> <p>Depending on your environment, decryption and decompression on the client may provide better performance.</p> <p>Default value: RESTORE_DECRYPT_LOCAL = 0</p> <p>Possible values: 0 enables decryption and decompression on the media server; 1 enables decryption and decompression on the client.</p>
SEGKSIZE	<p>The default file segment size in kilobytes.</p> <p>Default value: SEGKSIZE = 128</p> <p>Possible values: 32 to 16384 KBs, increments of 32 only</p> <p>Warning: Changing this value may reduce capacity and decrease performance. Change this value only when directed to do so by a Symantec representative.</p> <p>You can also specify the segment size for specific file types. See PREFERRED_EXT_SEGKSIZE.</p>
WS_RETRYCOUNT	<p>This parameter applies to the PureDisk Deduplication Option only. It does not affect NetBackup deduplication.</p> <p>Default value: WS_RETRYCOUNT = 3</p> <p>Possible values: Integer</p>
WS_TIMEOUT	<p>This parameter applies to the PureDisk Deduplication Option only. It does not affect NetBackup deduplication.</p> <p>Default value: WS_TIMEOUT = 120</p> <p>Possible values: Integer</p>

Editing the pd.conf deduplication file

If you change the pd.conf file on a host, it changes the settings for that host only. If you want the same settings for all of the hosts that deduplicate data, you must change the pd.conf file on all of the hosts.

See “[About the `pd.conf` configuration file for NetBackup deduplication](#)” on page 119.

See “[pd.conf file settings for NetBackup deduplication](#) ” on page 119.

To edit the `pd.conf` file

- 1 Use a text editor to open the `pd.conf` file.

The `pd.conf` file resides in the following directories:

- (UNIX) `/usr/opensv/lib/ost-plugins/`
 - (Windows) `install_path\Veritas\NetBackup\bin\ost-plugins`
- 2 To activate a setting, remove the pound character (#) in column 1 from each line that you want to edit.
 - 3 To change a setting, specify a new value.

Note: The spaces to the left and right of the equal sign (=) in the file are significant. Ensure that the space characters appear in the file after you edit the file.

- 4 Save and close the file.
- 5 Restart the NetBackup Remote Manager and Monitor Service (`nbrmms`) on the host.

About the `contentrouter.cfg` file for NetBackup deduplication

The `contentrouter.cfg` file contains various configuration settings that control some of the operations of your deduplication environment.

Usually, you do not need to change settings in the file. However, in some cases, you may be directed to change settings by a Symantec support representative.

The `contentrouter.cfg` file resides in the following directories:

- (UNIX) `storage_path/etc/puredisk`
- (Windows) `storage_path\etc\puredisk`

The NetBackup documentation exposes only some of the `contentrouter.cfg` file parameters. Those parameters appear in topics that describe a task or process to change configuration settings.

Note: Change values in the `contentrouter.cfg` only when directed to do so by the NetBackup documentation or by a Symantec representative.

About saving the deduplication storage server configuration

You can save your storage server settings in a text file. A saved storage server configuration file contains the configuration settings for your storage server. It also contains status information about the storage. A saved configuration file may help you with recovery of your storage server. Therefore, Symantec recommends that you get the storage server configuration and save it in a file. The file does not exist unless you create it.

The following is an example of a populated configuration file:

```
V7.0 "storagepath" "D:\DedupeStorage" string
V7.0 "spalogpath" "D:\DedupeStorage\log" string
V7.0 "dbpath" "D:\DedupeStorage" string
V7.0 "required_interface" "HOSTNAME" string
V7.0 "spalogretention" "7" int
V7.0 "verboselevel" "3" int
V7.0 "replication_target(s)" "none" string
V7.0 "Storage Pool Size" "698.4GB" string
V7.0 "Storage Pool Used Space" "132.4GB" string
V7.0 "Storage Pool Available Space" "566.0GB" string
V7.0 "Catalog Logical Size" "287.3GB" string
V7.0 "Catalog files Count" "1288" string
V7.0 "Space Used Within Containers" "142.3GB" string
```

V7.0 represents the version of the I/O format not the NetBackup release level. The version may differ on your system.

If you get the storage server configuration when the server is not configured or is down and unavailable, NetBackup creates a template file. The following is an example of a template configuration file:

```
V7.0 "storagepath" " " string
V7.0 "spalogin" " " string
V7.0 "spasswd" " " string
V7.0 "spalogretention" "7" int
V7.0 "verboselevel" "3" int
V7.0 "dbpath" " " string
V7.0 "required_interface" " " string
```


To use a storage server configuration file for recovery, you must edit the file so that it includes only the information that is required for recovery.

See “[Saving the deduplication storage server configuration](#)” on page 129.

See “[Editing a deduplication storage server configuration file](#)” on page 129.

See “[Setting the deduplication storage server configuration](#)” on page 131.

Saving the deduplication storage server configuration

Symantec recommends that you save the storage server configuration in a file. A storage server configuration file can help with recovery.

See “[About saving the deduplication storage server configuration](#)” on page 128.

See “[Recovering from a deduplication storage server disk failure](#)” on page 209.

See “[Recovering from a deduplication storage server failure](#)” on page 211.

To save the storage server configuration

◆ On the master server, enter the following command:

```
UNIX: /usr/opensv/netbackup/bin/admincmd/nbdevconfig -getconfig  
-storage_server sshostname -stype PureDisk -configlist file.txt
```

```
Windows: install_path\NetBackup\bin\admincmd\nbdevconfig -getconfig  
-storage_server sshostname -stype PureDisk -configlist file.txt
```

For *sshostname*, use the name of the storage server. For *file.txt*, use a file name that indicates its purpose.

If you get the file when a storage server is not configured or is down and unavailable, NetBackup creates a template file.

Editing a deduplication storage server configuration file

To use a storage server configuration file for recovery, it must contain only the required information. You must remove any point-in-time status information. (Status information is only in a configuration file that was saved on an active storage server.) You also must add several configuration settings that are not included in a saved configuration file or a template configuration file.

[Table 5-12](#) shows the configuration lines that are required.

Table 5-12 Required lines for a recovery file

Configuration setting	Description
V7.0 "storagepath" " " string	The value should be the same as the value that was used when you configured the storage server.
V7.0 "spalogpath" " " string	For the <code>spalogpath</code> , use the <code>storagepath</code> value and append <code>log</code> to the path. For example, if the <code>storagepath</code> is <code>D:\DedupeStorage</code> , enter <code>D:\DedupeStorage\log</code> .
V7.0 "dbpath" " " string	If the database path is the same as the <code>storagepath</code> value, enter the same value for <code>dbpath</code> . Otherwise, enter the path to the database.
V7.0 "required_interface" " " string	A value for <code>required_interface</code> is required only if you configured one initially; if a specific interface is not required, leave it blank. In a saved configuration file, the required interface defaults to the computer's hostname.
V7.0 "spalogretention" "7" int	Do not change this value.
V7.0 "verboselevel" "3" int	Do not change this value.
V7.0 "replication_target(s)" "none" string	A value for <code>replication_target(s)</code> is required only if you configured optimized duplication. Otherwise, do not edit this line.
V7.0 "spalogin" "username" string	Replace <code>username</code> with the NetBackup Deduplication Engine user ID.
V7.0 "spapasswd" "password" string	Replace <code>password</code> with the password for the NetBackup Deduplication Engine user ID.

See [“About saving the deduplication storage server configuration”](#) on page 128.

See [“Recovering from a deduplication storage server disk failure”](#) on page 209.

See [“Recovering from a deduplication storage server failure”](#) on page 211.

To edit the storage server configuration

- 1 If you did not save a storage server configuration file, get a storage server configuration file.

See [“Saving the deduplication storage server configuration”](#) on page 129.

- 2 Use a text editor to enter, change, or remove values.

Remove lines from and add lines to your file until only the required lines (see [Table 5-12](#)) are in the configuration file. Enter or change the values between the second set of quotation marks in each line. A template configuration file has a space character (" ") between the second set of quotation marks.

Setting the deduplication storage server configuration

You can set the storage server configuration (that is, configure the storage server) by importing the configuration from a file. Setting the configuration can help you with recovery of your environment.

See [“Recovering from a deduplication storage server disk failure”](#) on page 209.

See [“Recovering from a deduplication storage server failure”](#) on page 211.

To set the configuration, you must have an edited storage server configuration file.

See [“About saving the deduplication storage server configuration”](#) on page 128.

See [“Saving the deduplication storage server configuration”](#) on page 129.

See [“Editing a deduplication storage server configuration file”](#) on page 129.

Note: The only time you should use the `nbdevconfig` command with the `-setconfig` option is for recovery of the host or the host disk.

To set the storage server configuration

- ◆ On the master server, run the following command:

```
UNIX: /usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig  
-storage_server sshostname -stype PureDisk -configlist file.txt
```

```
Windows: install_path\NetBackup\bin\admincmd\nbdevconfig -setconfig  
-storage_server sshostname -stype PureDisk -configlist file.txt
```

For `sshostname`, use the name of the storage server. For `file.txt`, use the name of the file that contains the configuration.

About the deduplication host configuration file

Each NetBackup host that is used for deduplication has a configuration file; the file name matches the name of the storage server, as follows:

storage_server_name.cfg

The *storage_server_name* is the fully qualified domain name if that was used to configure the storage server. For example, if the storage server name is

DedupeServer.symantecs.org, the configuration file name is

DedupeServer.symantecs.org.cfg.

The following is the location of the file:

UNIX: `/usr/opensv/lib/ost-plugins`

Windows: `install_path\Veritas\NetBackup\bin\ost-plugins`

Deleting a deduplication host configuration file

You may need to delete the configuration file from the deduplication hosts. For example, to reconfigure your deduplication environment or disaster recovery may require that you delete the configuration file on the servers on which it exists.

See [“About the deduplication host configuration file”](#) on page 132.

To delete the host configuration file

- ◆ Delete the file; its location depends on the operating system type, as follows:

UNIX: `/usr/opensv/lib/ost-plugins`

Windows: `install_path\Veritas\NetBackup\bin\ost-plugins`

The following is an example of the host configuration file name of a server that has a fully qualified domain name:

DedupeServer.symantecs.org.cfg

Resetting the deduplication registry

If you reconfigure your deduplication environment, one of the steps is to reset the deduplication registry.

See [“Changing the deduplication storage server name or storage path”](#) on page 155.

Warning: Only follow these procedures if you are reconfiguring your storage server and storage paths.

The procedure differs on UNIX and on Windows.

To reset the deduplication registry file on UNIX and Linux

- ◆ Enter the following commands on the storage server to reset the deduplication registry file:

```
rm /etc/pdregistry.cfg
cp -f /usr/opensv/pdde/pdconfigure/cfg/userconfigs/pdregistry.cfg
   /etc/pdregistry.cfg
```

To reset the deduplication registry on Windows

- 1 Delete the contents of the following keys in the Windows registry:

- HKLM\SOFTWARE\Symantec\PureDisk\Agent\ConfigFilePath
- HKLM\SOFTWARE\Symantec\PureDisk\Agent\EtcPath

Warning: Editing the Windows registry may cause unforeseen results.

- 2 Delete the storage path in the following key in the Windows key. That is, delete everything after `postgresql-8.3 -D` in the key.

```
HKLM\SYSTEM\ControlSet001\Services\postgresql-8.3\ImagePath
```

For example, in the following example registry key, you would delete the content of the key that is in italic type:

```
"C:\Program Files\Veritas\pdde\pddb\bin\pg_ctl.exe" runservice
-N postgresql-8.3 -D "D:\DedupeStorage\databases\pddb\data" -w
```

The result is as follows:

```
"C:\Program Files\Veritas\pdde\pddb\bin\pg_ctl.exe" runservice
-N postgresql-8.3 -D
```

Configuring deduplication log file timestamps on Windows

The default configuration for the PostgreSQL database does not add timestamps to log entries on Windows systems. Therefore, Symantec recommends that you edit the PostgreSQL configuration file on Windows hosts so timestamps are added to the log file.

To configure log file timestamps on Windows

- 1 Use a text editor to open the following file:

```
dbpath\databases\pddb\data\postgresql.conf
```

The database path may be the same as the configured storage path.

- 2 In the line that begins with `log_line_prefix`, change the value from `%t` to `%t`. (That is, remove one of the percent signs (%).)
- 3 Save the file.
- 4 Run the following command:

```
install_path\Veritas\pdde\pddb\bin\pg_ctl reload -D  
dbpath\databases\pddb\data
```

If the command output does not include `server signaled`, use Windows Computer Management to restart the PostgreSQL Server 8.3 service.

See [“About deduplication logs”](#) on page 187.

Setting NetBackup configuration options by using `bpsetconfig`

Symantec recommends that you use the **NetBackup Administration Console Host Properties** to configure NetBackup.

However, some configuration options cannot be set by using the **Administration Console**. You can set those configuration options by using the `bpsetconfig` command. Configuration options are key and value pairs, as shown in the following examples:

- `SERVER = server1.symantecs.org`
- `CLIENT_READ_TIMEOUT = 300`
- `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE`

Alternatively, on UNIX systems you can set configuration options in the `bp.conf` file.

See the *NetBackup Administrator's Guide for UNIX and Linux, Volume I*.

To set configuration options by using the `bpsetconfig` command

- 1** On the host on which you want to set configuration options, invoke the `bpsetconfig` command, as follows:

UNIX: `/usr/opensv/netbackup/bin/admincmd/bpsetconfig`

Windows: `install_path\NetBackup\bin\admincmd\bpsetconfig`

- 2** At the `bpsetconfig` prompt, enter the key and the value pairs of the configuration options that you want to set, one pair per line.

Ensure that you understand the values that are allowed and the format of any new options that you add.

You can change existing key and value pairs.

You can add key and value pairs.

- 3** To save the configuration changes, type the following, depending on the operating system:

UNIX: `Ctrl + D Enter`

Windows: `Ctrl + Z Enter`

Monitoring deduplication activity

This chapter includes the following topics:

- [Monitoring the deduplication rate](#)
- [Viewing deduplication job details](#)
- [About deduplication storage capacity and usage reporting](#)
- [About deduplication container files](#)
- [Viewing storage usage within deduplication container files](#)
- [Viewing disk reports](#)
- [Monitoring deduplication processes](#)
- [Reporting on Auto Image Replication jobs](#)

Monitoring the deduplication rate

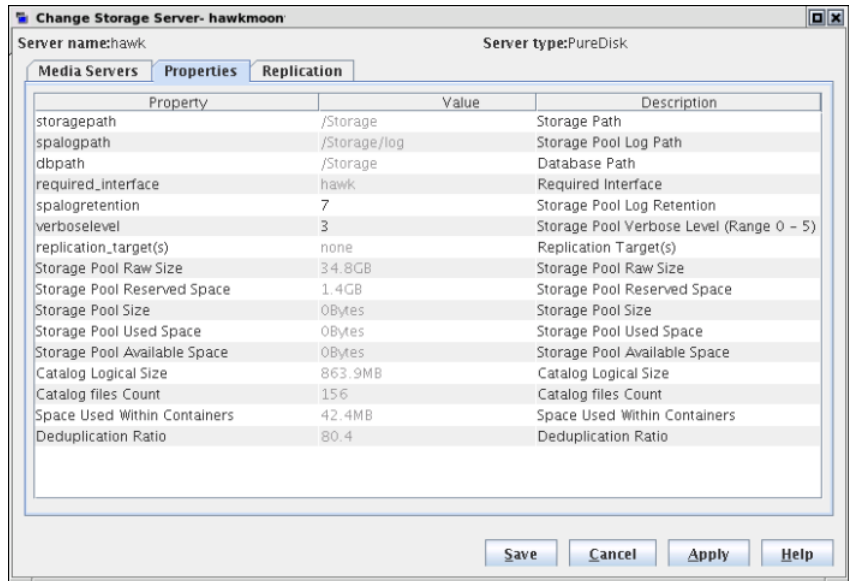
The deduplication rate is the percentage of data that was stored already. That data is not stored again.

The following methods show the deduplication rate:

- [To view the global deduplication ratio](#)
- [To view the deduplication rate for a backup job in the Activity Monitor](#)

To view the global deduplication ratio

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Credentials > Storage Server**
- 2 Select the deduplication storage server.
- 3 On the **Edit** menu, select **Change**.
- 4 In the **Change Storage Server** dialog box, select the **Properties** tab. The **Deduplication Ratio** field displays the ratio.



To view the deduplication rate for a backup job in the Activity Monitor

- 1 In the NetBackup Administration Console, click **Activity Monitor**.
- 2 Click the **Jobs** tab.

The **Deduplication Rate** column shows the rate for each job.

Many factors affect deduplication performance.

See [“About deduplication performance”](#) on page 52.

Viewing deduplication job details

Use the NetBackup Activity Monitor to view deduplication job details.

To view deduplication job details

- 1 In the **NetBackup Administration Console**, click **Activity Monitor**.
- 2 Click the **Jobs** tab.
- 3 To view the details for a specific job, double-click on the job that is displayed in the **Jobs** tab pane.
- 4 In the **Job Details** dialog box, click the **Detailed Status** tab.

The deduplication job details are described in a different topic.

See [“Deduplication job details”](#) on page 139.

Deduplication job details

In the **Job Details** dialog box, the details of a deduplication job depend on whether the job is media server deduplication or client-side deduplication. [Table 6-1](#) describes the job types.

Table 6-1 Deduplication job type and description

Job type	Description
Media server deduplication job details	<p>For media server deduplication, the Detailed Status tab shows the deduplication rate on the server that performed the deduplication. The following job details excerpt shows details for a client for which Server_A deduplicated the data (the dedup field shows the deduplication rate):</p> <pre>10/6/2010 10:02:09 AM - Info Server_A(pid=30695) StorageServer=PureDisk:Server_A; Report=PDDO Stats for (Server_A): scanned: 30126998 KB, stream rate: 162.54 MB/sec, CR sent: 1720293 KB, dedup: 94.3%, cache hits: 214717 (94.0%)</pre>

Table 6-1 Deduplication job type and description (*continued*)

Job type	Description
Client-side deduplication job details	<p>For client-side deduplication jobs, the Detailed Status tab shows two deduplication rates. The first deduplication rate is always for the client data. The second deduplication rate is for the metadata (disk image header and True Image Restore information (if applicable)). That information is always deduplicated on a server; typically, deduplication rates for that information are zero or very low. The following job details example excerpt shows the two rates. The 10/8/2009 11:58:09 PM entry is for the client data; the 10/8/2010 11:58:19 PM entry is for the metadata.</p> <pre> 10/8/2010 11:54:21 PM - Info Server_A(pid=2220) Using OpenStorage client direct to backup from client Client_B to Server_A 10/8/2009 11:58:09 PM - Info Server_A(pid=2220) StorageServer=PureDisk:Server_A; Report=PDDO Stats for (Server_A: scanned: 3423425 KB, stream rate: 200.77 MB/sec, CR sent: 122280 KB, dedup: 96.4%, cache hits: 49672 (98.2%) 10/8/2010 11:58:09 PM - Info Server_A(pid=2220) Using the media server to write NBU data for backup Client_B_1254987197 to Server_A 10/8/2010 11:58:19 PM - Info Server_A(pid=2220) StorageServer=PureDisk:Server_A; Report=PDDO Stats for (Server_A: scanned: 17161 KB, stream rate: 1047.42 MB/sec, CR sent: 17170 KB, dedup: 0.0%, cache hits: 0 (0.0%) the requested operation was successfully completed(0) </pre>

[Table 6-2](#) describes the deduplication activity fields.

Table 6-2 Deduplication activity field descriptions

Field	Description
cache hits	<p>The percentage of time that the local fingerprint cache contained a record of the segment. The deduplication plug-in did not have to query the database about the segment.</p> <p>If the <code>pd.conf</code> file <code>FP_CACHE_LOCAL</code> parameter is set to 0 on the storage, the cache hits output is not included for the jobs that run on the storage server.</p> <p>See “pd.conf file settings for NetBackup deduplication” on page 119.</p>

Table 6-2 Deduplication activity field descriptions (*continued*)

Field	Description
CR sent	<p>The amount of data that is sent from the deduplication plug-in to the component that stores the data. In NetBackup, the NetBackup Deduplication Engine stores the data.</p> <p>If the storage server deduplicates the data, it does not travel over the network. The deduplicated data travels over the network when the deduplication plug-in runs on a computer other than the storage server, as follows:</p> <ul style="list-style-type: none"> ■ On a NetBackup client that deduplicates its own data (client-side deduplication). ■ On a fingerprinting media server that deduplicates the data. The plug-in on the fingerprinting server sends the data to the storage server, which writes it to a Media Server Deduplication Pool. ■ On a media server that then sends it to a PureDisk environment for storage. (In NetBackup, a PureDisk Storage Pool represents the storage of a PureDisk environment.)
CR sent over FC	The amount of data that is sent from the deduplication plug-in over Fibre Channel to the component that stores the data. In NetBackup, the NetBackup Deduplication Engine stores the data.
dedup	The percentage of data that was stored already. That data is not stored again.
scanned	The amount of data that the deduplication plug-in scanned.
stream rate	The speed of the scan: The kilobytes of data that are scanned divided by how long the scan takes.

About deduplication storage capacity and usage reporting

Several factors affect the expected NetBackup deduplication capacity and usage results, as follows:

- Expired backups may not change the available size and the used size. An expired backup may have no unique data segments. Therefore, the segments remain valid for other backups.
- NetBackup Deduplication Manager clean-up may not have run yet. The Deduplication Manager performs clean up twice a day. Until it performs clean-up, deleted image fragments remain on disk.

If you use operating system tools to examine storage space usage, their results may differ from the usage reported by NetBackup, as follows:

- NetBackup usage data includes reserved space that the operating system tools do not include.
- If other applications use the storage, NetBackup cannot report usage accurately. NetBackup requires exclusive use of the storage.

[Table 6-3](#) describes the options for monitoring capacity and usage.

Table 6-3 Capacity and usage reporting

Option	Description
Change Storage Server dialog box	<p>The Change Storage Server dialog box Properties tab displays storage capacity and usage. It also displays the global deduplication ratio.</p> <p>This dialog box displays the most current capacity usage that is available in the NetBackup Administration Console.</p> <p>You can see an example of the dialog box in a different topic. See “Monitoring the deduplication rate” on page 137.</p>
Disk Pools window	<p>The Disk Pools window of the Administration Console displays values that were stored when NetBackup polled the disk pools. NetBackup polls every 5 minutes; therefore, the value may not be as current as the value that is displayed in the Change Storage Server dialog box.</p> <p>To display the window, expand Media and Device Management > Devices > Disk Pools.</p>
View container command	<p>A command that is installed with NetBackup provides a view of storage capacity and usage within the deduplication container files.</p> <p>See “About deduplication container files” on page 143.</p> <p>See “Viewing storage usage within deduplication container files” on page 143.</p>
License Keys dialog box	<p>The summary of active capacity-based license features in the NetBackup License Keys dialog box. The summary displays the storage capacity for which you are licensed and the capacity used. It does not display the amount of physical storage space.</p> <p>On the Help menu in the NetBackup Administration Console, select License Keys.</p>

Table 6-3 Capacity and usage reporting (*continued*)

Option	Description
Disk Pool Status report	The Disk Pool Status report displays the state of the disk pool and usage information. See “Viewing disk reports” on page 145.
Disk Logs report	The Disk Logs report displays event and message information. A useful event for monitoring capacity is event 1044; the following is the description of the event in the Disk Logs report: The usage of one or more system resources has exceeded a warning level. The threshold for this message is at 96% capacity. No more data can be stored. See “Viewing disk reports” on page 145. See “Deduplication event codes and messages” on page 202.
The <code>nbdevquery</code> command	The <code>nbdevquery</code> command shows the state of the disk volume and its properties and attributes. It also shows capacity, usage, and percent used. See “Determining the deduplication disk volume state” on page 171.
NetBackup OpsCenter	The NetBackup OpsCenter also provides information about storage capacity and usage. See the <i>NetBackup OpsCenter Administrator's Guide</i> .

About deduplication container files

The deduplication storage implementation allocates container files to hold backup data. Deleted segments can leave free space in containers files, but the container file sizes do not change. Segments are deleted from containers when backup images expire and the NetBackup Deduplication Manager performs clean-up.

Viewing storage usage within deduplication container files

A deduplication command reports on storage usage within containers.

The following is the pathname of the command:

- **UNIX and Linux:** `/usr/opensv/pdde/pdcr/bin/crcontrol.`
- **Windows:** `install_path\Veritas\pdde\Crcontrol.exe.`

To view storage usage within deduplication container files

- ◆ Use the `crcontrol` command and the `--dsstat` option on the deduplication storage server. (For help with the command options, use the `--help` option.)

The following is an example of the command usage on a Windows deduplication storage server.

```
C:\Program Files\Veritas\pdde>Crcontrol.exe --dsstat

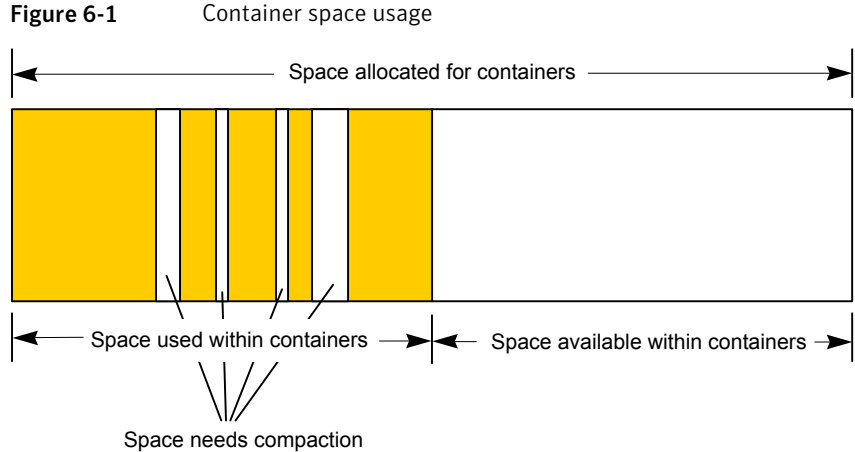
***** Data Store statistics *****
Data storage      Raw    Size   Used   Avail  Use%
                  62.8T 60.2T 58.7T  1.5T  98%

Number of containers      : 239731
Average container size    : 268120188 bytes (255.70MB)
Space allocated for containers : 64276720994658 bytes (58.46TB)
Space used within containers : 63904602067295 bytes (58.12TB)
Space available within containers: 372118927363 bytes (346.56GB)
Space needs compaction    : 154014169 bytes (146.88MB)
Reserved space           : 2865754513408 bytes (2.61TB)
Reserved space percentage : 4.2%
Records marked for compaction : 29620
Active records           : 912115621
Total records            : 912145241
```

From the command output, you can determine the following:

Raw	The raw size of the storage before the space reserved for meta data.
Size	Raw minus Reserved space.
Used	The file system used space minus Space available within containers minus Space needs compaction. NetBackup obtains the file system used space from the operating system.
Avail	Size minus Used.
Use%	Used divided by Size.

Figure 6-1 is a visual representation of the deduplication storage space.



The NetBackup Deduplication Manager checks the storage space every 20 seconds. It then periodically compacts the space available inside the container files. Therefore, space within a container is not available as soon as it is free. Various internal parameters control whether a container file is compacted. Although space may be available within a container file, the file may not be eligible for compaction.

Viewing disk reports

The NetBackup disk reports include information about the disk pools, disk storage units, disk logs, images that are stored on disk media, and storage capacity.

[Table 6-4](#) describes the disk reports available.

Table 6-4 Disk reports

Report	Description
Images on Disk	<p>The Images on Disk report generates the image list present on the disk storage units that are connected to the media server. The report is a subset of the Images on Media report; it shows only disk-specific columns.</p> <p>The report provides a summary of the storage unit contents. If a disk becomes bad or if a media server crashes, this report can let you know what data is lost.</p>

Table 6-4 Disk reports (*continued*)

Report	Description
Disk Logs	<p>The Disk Logs report displays the media errors or the informational messages that are recorded in the NetBackup error catalog. The report is a subset of the Media Logs report; it shows only disk-specific columns.</p> <p>The report also includes information about deduplicated data integrity checking.</p> <p>See “About deduplication data integrity checking” on page 175.</p> <p>Either PureDisk or Symantec Deduplication Engine in the description identifies a deduplication message. The identifiers are generic because the deduplication engine does not know which application consumes its resources. NetBackup, Symantec Backup Exec, and NetBackup PureDisk are Symantec applications that use deduplication.</p>
Disk Storage Unit	<p>The Disk Storage Unit Status report displays the state of disk storage units in the current NetBackup configuration.</p> <p>For disk pool capacity, see the disk pools window in Media and Device Management > Devices > Disk Pools.</p> <p>Multiple storage units can point to the same disk pool. When the report query is by storage unit, the report counts the capacity of disk pool storage multiple times.</p>
Disk Pool Status	<p>The Disk Pool Status report displays the state of disk pool and usage information.</p>

To view disk reports

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Reports > Disk Reports**.
- 2 Select the name of a disk report.
- 3 In the right pane, select the report settings.
- 4 Click **Run Report**.

Monitoring deduplication processes

The following table shows the deduplication processes about which NetBackup reports:

See [“Deduplication storage server components”](#) on page 215.

Table 6-5 Where to monitor the main deduplication processes

What	Where to monitor it
NetBackup Deduplication Engine	<p>On Windows systems, in the NetBackup Administration Console Activity Monitor Services tab.</p> <p>On UNIX, the NetBackup Deduplication Engine appears as <code>spoold</code> in the Administration Console Activity Monitor Daemons tab.</p> <p>The NetBackup <code>bpps</code> command shows the <code>spoold</code> process.</p>
NetBackup Deduplication Manager	<p>On Windows systems, NetBackup Deduplication Manager in the Activity Monitor Services tab.</p> <p>On UNIX, the NetBackup Deduplication Manager appears as <code>spad</code> in the Administration Console Activity Monitor Daemons tab.</p> <p>The NetBackup <code>bpps</code> command shows the <code>spad</code> process.</p>
The database processes (<code>postgres</code>)	<p>On Windows systems, the <code>postgres</code> database processes appear in the Activity Monitor Processes tab.</p> <p>The NetBackup <code>bpps</code> command shows the <code>postgres</code> processes.</p>

Reporting on Auto Image Replication jobs

The Activity Monitor displays both the **Replication** job and the **Import** job in a configuration that replicates to a target master server domain.

Table 6-6 Auto Image Replication jobs in the Activity Monitor

Job type	Description
Replication	<p>The job that replicates a backup image to a target master displays in the Activity Monitor as a Replication job. The Target Master label displays in the Storage Unit column for this type of job.</p> <p>Similar to other Replication jobs, the job that replicates images to a target master can work on multiple backup images in one instance.</p> <p>The detailed status for this job contains a list of the backup IDs that were replicated.</p>

Table 6-6 Auto Image Replication jobs in the Activity Monitor *(continued)*

Job type	Description
Import	<p>The job that imports a backup copy into the target master domain displays in the Activity Monitor as an Import job. An Import job can import multiple copies in one instance. The detailed status for an Import job contains a list of processed backup IDs and a list of failed backup IDs.</p> <p>Note that a successful replication does not confirm that the image was imported at the target master.</p> <p>If the SLP names or data classifications are not the same in both domains, the Import job fails and NetBackup does not attempt to import the image again.</p> <p>Failed Import jobs fail with a status 191 and appear in the Problems report when run on the target master server.</p> <p>The image is expired and deleted during an Image Cleanup job. Note that the originating domain (Domain 1) does not track failed imports.</p>

Managing deduplication

This chapter includes the following topics:

- [Managing deduplication servers](#)
- [Managing NetBackup Deduplication Engine credentials](#)
- [Managing deduplication disk pools](#)
- [Deleting backup images](#)
- [Disabling client-side deduplication for a client](#)
- [About deduplication queue processing](#)
- [Processing the deduplication transaction queue manually](#)
- [About deduplication data integrity checking](#)
- [Configuring deduplication data integrity checking behavior](#)
- [About managing storage read performance](#)
- [About deduplication storage rebasing](#)
- [Resizing the deduplication storage partition](#)
- [About restoring files at a remote site](#)
- [About restoring from a backup at a target master domain](#)
- [Specifying the restore server](#)

Managing deduplication servers

After you configure deduplication, you can perform various tasks to manage deduplication servers.

See [“Viewing deduplication storage servers”](#) on page 150.

See [“Determining the deduplication storage server state”](#) on page 150.

See [“Viewing deduplication storage server attributes”](#) on page 151.

See [“Setting deduplication storage server attributes”](#) on page 152.

See [“Changing deduplication storage server properties”](#) on page 153.

See [“Clearing deduplication storage server attributes”](#) on page 154.

See [“About changing the deduplication storage server name or storage path”](#) on page 155.

See [“Changing the deduplication storage server name or storage path”](#) on page 155.

See [“Removing a load balancing server”](#) on page 157.

See [“Deleting a deduplication storage server”](#) on page 158.

See [“Deleting the deduplication storage server configuration”](#) on page 159.

Viewing deduplication storage servers

Use the NetBackup Administration Console to view a list of deduplication storage servers already configured.

To view deduplication storage servers

- ◆ In the NetBackup Administration Console, expand **Media and Device Management > Credentials > Storage Server**.

The **All Storage Servers** pane shows all configured deduplication storage servers. deduplication storage servers show **PureDisk** in the **Disk Type** column.

Determining the deduplication storage server state

Use the NetBackup `nbdevquery` command to determine the state of a deduplication storage server. The state is either UP or DOWN.

To determine deduplication storage server state

- ◆ Run the following command on the NetBackup master server or a deduplication storage server:

```
UNIX: /usr/opensv/netbackup/bin/admincmd/nbdevquery -liststs  
-storage_server server_name -stype PureDisk -U
```

```
Windows: install_path\NetBackup\bin\admincmd\nbdevquery -liststs  
-storage_server server_name -stype PureDisk -U
```

The following is example output:

```
Storage Server      : bit  
Storage Server Type : PureDisk  
Storage Type       : Formatted Disk, Network Attached  
State              : UP
```

This example output is shortened; more flags may appear in actual output.

Viewing deduplication storage server attributes

Use the NetBackup `nbdevquery` command to view the deduplication storage server attributes.

The *server_name* you use in the `nbdevquery` command must match the configured name of the storage server. If the storage server name is its fully-qualified domain name, you must use that for *server_name*.

To view deduplication storage server attributes

- ◆ The following is the command syntax to set a storage server attribute. Run the command on the NetBackup master server or on the deduplication storage server:

```
UNIX: /usr/opensv/netbackup/bin/admincmd/nbdevquery -liststs  
-storage_server server_name -stype PureDisk -U
```

```
Windows: install_path\NetBackup\bin\admincmd\nbdevquery -liststs  
-storage_server server_name -stype PureDisk -U
```

The following is example output:

```
Storage Server      : bit  
Storage Server Type : PureDisk  
Storage Type       : Formatted Disk, Network Attached  
State              : UP  
Flag               : OpenStorage  
Flag               : CopyExtents  
Flag               : AdminUp  
Flag               : InternalUp  
Flag               : LifeCycle  
Flag               : CapacityMgmt  
Flag               : OptimizedImage  
Flag               : FT-Transfer
```

This example output is shortened; more flags may appear in actual output.

Setting deduplication storage server attributes

You may have to set storage server attributes to enable new functionality.

If you set an attribute on the storage server, you may have to set the same attribute on existing deduplication pools. The overview or configuration procedure for the new functionality describes the requirements.

See “[Setting deduplication disk pool attributes](#)” on page 164.

To set a deduplication storage server attribute

- 1 The following is the command syntax to set a storage server attribute. Run the command on the master server or on the storage server.

```
nbdevconfig -changests -storage_server storage_server -stype  
PureDisk -setattribute attribute
```

The following describes the options that require the arguments that are specific to your domain:

<code>-storage_server</code> <code>storage_server</code>	The name of the storage server.
<code>-setattribute</code> <code>attribute</code>	The <i>attribute</i> is the name of the argument that represents the new functionality. For example, OptimizedImage specifies that the environment supports the optimized synthetic backup method.

The following is the path to the `nbdevconfig` command:

- UNIX: `/usr/opensv/netbackup/bin/admincmd`
- Windows: `install_path\NetBackup\bin\admincmd`

- 2 To verify, view the storage server attributes.

See [“Viewing deduplication storage server attributes”](#) on page 151.

See [“About optimized synthetic backups and deduplication”](#) on page 35.

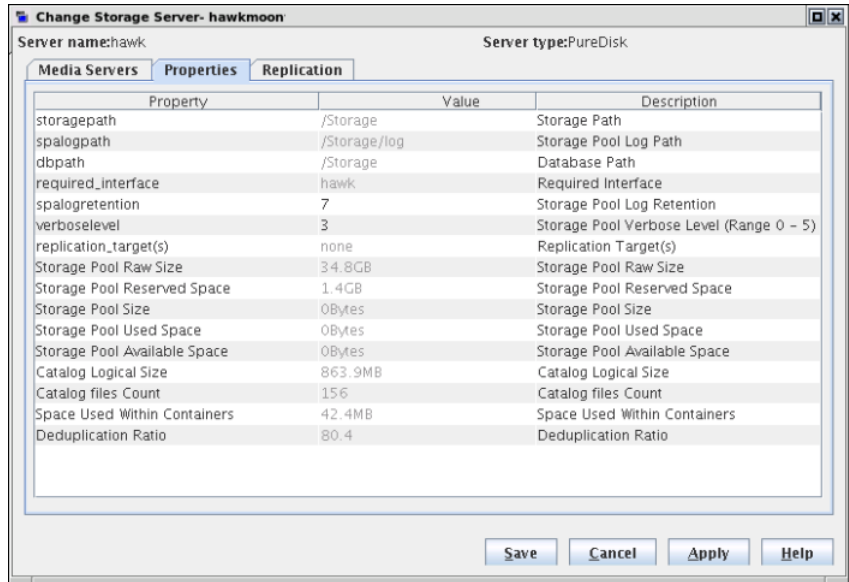
Changing deduplication storage server properties

You can change the retention period and logging level for the NetBackup Deduplication Manager.

To change deduplication storage server properties

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Credentials > Storage Server**
- 2 Select the deduplication storage server.
- 3 On the **Edit** menu, select **Change**.

- 4 In the **Change Storage Server** dialog box, select the **Properties** tab.



- 5 For the property to change, select the value in the **Value** column.
- 6 Change the value.
- 7 Click **OK**.

Clearing deduplication storage server attributes

Use the `nbdevconfig` command to remove storage server attributes.

To clear deduplication storage server attributes

- ◆ Run the following command on the NetBackup master server or on a storage server:

```
nbdevconfig -changests -storage_server storage_server -stype  
PureDisk -clearattribute attribute
```

`-storage_server` The name of the storage server.
`storage_server`

`-setattribute` The *attribute* is the name of the argument that represents
`attribute` the functionality.

The following is the path to the `nbdevconfig` command:

- UNIX: `/usr/opensv/netbackup/bin/admincmd`
- Windows: `install_path\NetBackup\bin\admincmd`

About changing the deduplication storage server name or storage path

You can change the storage server host name and the storage path of an existing NetBackup deduplication environment.

The following are several use cases that require changing an existing deduplication environment:

- You want to change the host name. For example, the name of host A was changed to B or a new network card was installed with a private interface C. To use the host name B or the private interface C, you must reconfigure the storage server.
See [“Changing the deduplication storage server name or storage path”](#) on page 155.
- You want to change the storage path. To do so, you must reconfigure the storage server with the new path.
See [“Changing the deduplication storage server name or storage path”](#) on page 155.
- You need to reuse the storage for disaster recovery. The storage is intact, but the storage server was destroyed. To recover, you must configure a new storage server.
In this scenario, you can use the same host name and storage path or use different ones.

Changing the deduplication storage server name or storage path

Two aspects of a NetBackup deduplication configuration exist: the record of the deduplication storage in the EMM database and the physical presence of the storage on disk (the populated storage directory).

Warning: Deleting valid backup images may cause data loss.

See [“About changing the deduplication storage server name or storage path”](#) on page 155.

Table 7-1 Changing the storage server name or storage path

Step	Task	Procedure
Step 1	Ensure that no deduplication activity occurs	Deactivate all backup policies that use deduplication storage. See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .
Step 2	Expire the backup images	Expire all backup images that reside on the deduplication disk storage. Warning: Do not delete the images. They are imported back into NetBackup later in this process. If you use the <code>bpexpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter. See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .
Step 3	Delete the storage units that use the disk pool	See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .
Step 4	Delete the disk pool	See “Deleting a deduplication disk pool” on page 172.
Step 5	Delete the deduplication storage server	See “Deleting a deduplication storage server” on page 158.
Step 6	Delete the configuration	Delete the deduplication configuration. See “Deleting the deduplication storage server configuration” on page 159.
Step 7	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See “Deleting a deduplication host configuration file” on page 132.
Step 8	Change the storage server name or the storage location	See the computer or the storage vendor's documentation. See “About the deduplication storage paths” on page 69.
Step 9	Reconfigure the storage server	When you configure deduplication, select the host by the new name and enter the new storage path (if you changed the path). You can also use a new network interface.
Step 10	Import the backup images	See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .

Removing a load balancing server

You can remove a load balancing server from a deduplication node. The media server no longer deduplicates client data.

See “[About NetBackup deduplication servers](#)” on page 25.

After you remove the load balancing server, restart the NetBackup Enterprise Media Manager service. The NetBackup disk polling service may try to use the removed server to query for disk status. Because the server is no longer a load balancing server, it cannot query the disk storage. Consequently, NetBackup may mark the disk volume as DOWN. When the EMM service restarts, it chooses a different deduplication server to monitor the disk storage.

If the host failed and is unavailable, you can use the `tpconfig` device configuration utility in menu mode to delete the server. However, you must run the `tpconfig` utility on a UNIX or Linux NetBackup server.

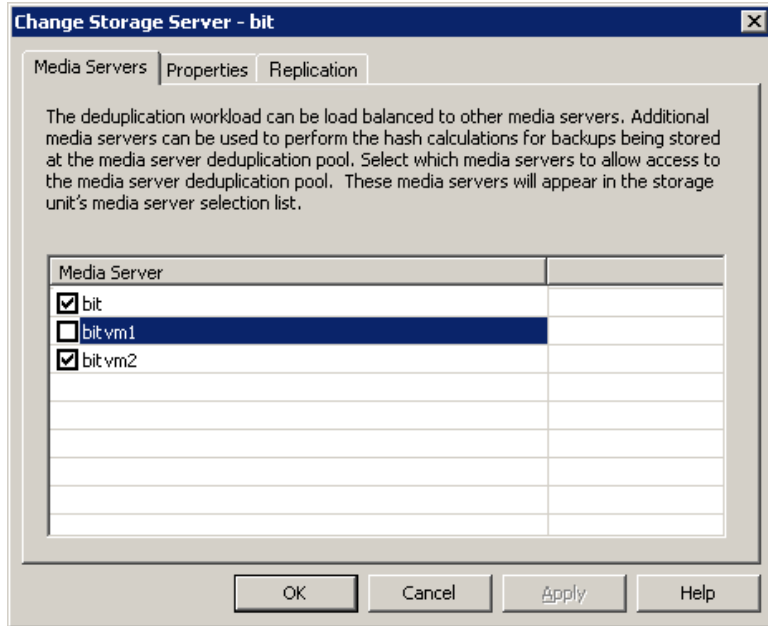
For procedures, see the *NetBackup Administrator's Guide for UNIX and Linux, Volume II*.

To remove a media server from a deduplication node

- 1 For every storage unit that specifies the media server in **Use one of the following media servers**, clear the check box that specifies the media server.

This step is not required if the storage unit is configured to use any available media server.
- 2 In the NetBackup Administration Console, expand **Media and Device Management > Credentials > Storage Server**.

- 3 Select the deduplication storage server, then select **Edit > Change**.



- 4 In the **Change Storage Server** dialog box, select the **Media Servers** tab.
- 5 Clear the check box of the media server you want to remove.
- 6 Click **OK**.

Deleting a deduplication storage server

If you delete a deduplication storage server, NetBackup deletes the host as a storage server and disables the deduplication storage server functionality on that media server.

NetBackup does not delete the media server from your configuration. To delete the media server, use the NetBackup `nbemmcmd` command.

Deleting the deduplication storage server does not alter the contents of the storage on physical disk. To protect against inadvertent data loss, NetBackup does not automatically delete the storage when you delete the storage server.

If a disk pool is configured from the disk volume that the deduplication storage server manages, you cannot delete the deduplication storage server.

Warning: Do not delete a deduplication storage server if its storage contains unexpired NetBackup images; if you do, data loss may occur.

To delete a deduplication storage server

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Credentials > Storage Server**
- 2 On the **Edit** menu, select **Delete**.
- 3 Click **Yes** in the confirmation dialog box.

See [“Changing the deduplication storage server name or storage path”](#) on page 155.

Deleting the deduplication storage server configuration

Use this procedure to delete a deduplication storage server configuration. The script used in this procedure deletes the active configuration and returns the configuration files to their installed, pre-configured state.

Only use this procedure when directed to from a process topic. A process topic is a high-level user task made up of a series of separate procedures.

See [“Changing the deduplication storage server name or storage path”](#) on page 155.

To delete the deduplication storage server configuration

- ◆ Delete the deduplication configuration by running one of the following scripts, depending on your operation system:

UNIX:

```
/usr/opensv/pdde/pdconfigure/scripts/installers/PDDE_deleteConfig.sh
```

Windows:install_path\Program

```
Files\Veritas\pdde\PDDE_deleteConfig.bat
```

About shared memory on Windows deduplication storage servers

On Windows deduplication servers, NetBackup uses shared memory for communication between the NetBackup Deduplication Manager (`spad.exe`) and the NetBackup Deduplication Engine (`spoold.exe`).

Usually, you should not be required to change the configuration settings for the shared memory functionality. However, after you upgrade to NetBackup 7.5, verify that the following shared memory values are set in the

`storage_path\etc\puredisk\agent.cfg` file:

```
SharedMemoryEnabled=1  
SharedMemoryBufferSize=262144  
SharedMemoryTimeout=3600
```

If the settings do not exist or their values differ from those in this topic, add or change them accordingly. Then, restart both the NetBackup Deduplication Manager (`spad.exe`) and the NetBackup Deduplication Engine (`spoold.exe`).

Managing NetBackup Deduplication Engine credentials

You can manage existing credentials in NetBackup.

See [“Determining which media servers have deduplication credentials”](#) on page 160.

See [“Adding NetBackup Deduplication Engine credentials”](#) on page 160.

See [“Changing NetBackup Deduplication Engine credentials”](#) on page 161.

See [“Deleting credentials from a load balancing server”](#) on page 161.

Determining which media servers have deduplication credentials

You can determine which media servers have credentials configured for the NetBackup Deduplication Engine. The servers with credentials are load balancing servers.

To determine if NetBackup Deduplication Engine credentials exist

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Credentials > Storage Server**.
- 2 Select the storage server, then select **Edit > Change**.
- 3 In the **Change Storage Server** dialog box, select the **Media Servers** tab.

The media servers for which credentials are configured are checked.

Adding NetBackup Deduplication Engine credentials

You may need to add the NetBackup Deduplication Engine credentials to an existing storage server or load balancing server. For example, disaster recovery may require that you add the credentials.

Add the same credentials that you already use in your environment.

Another procedure exists to add a load balancing server to your configuration.

See [“Adding a deduplication load balancing server”](#) on page 117.

To add NetBackup Deduplication Engine credentials by using the tpconfig command

- ◆ On the host to which you want to add credentials, run the following command:

```
UNIX: /usr/opensv/volmgr/bin/tpconfig -add -storage_server  
sshhostname -stype PureDisk -sts_user_id UserID -password PassWord
```

```
Windows: install_path\Veritas\NetBackup\Volmgr\bin\tpconfig -add  
-storage_server sshhostname -stype PureDisk -sts_user_id UserID  
-password PassWord
```

For *sshhostname*, use the name of the storage server.

Changing NetBackup Deduplication Engine credentials

You cannot change the NetBackup Deduplication Engine credentials after you enter them. If you must change the credentials, contact your Symantec support representative.

See [“About NetBackup Deduplication Engine credentials”](#) on page 32.

Deleting credentials from a load balancing server

You may need to delete the NetBackup Deduplication Engine credentials from a load balancing server. For example, disaster recovery may require that you delete the credentials on a load balancing server.

Another procedure exists to remove a load balancing server from a deduplication node.

See [“Removing a load balancing server”](#) on page 157.

To delete credentials from a load balancing server

- ◆ On the load balancing server, run the following command:

```
UNIX: /usr/opensv/volmgr/bin/tpconfig -delete -storage_server  
sshhostname -stype PureDisk -sts_user_id UserID
```

```
Windows: install_path\Veritas\NetBackup\Volmgr\bin\tpconfig -delete  
-storage_server sshhostname -stype PureDisk -sts_user_id UserID
```

For *sshhostname*, use the name of the storage server.

Managing deduplication disk pools

After you configure NetBackup deduplication, you can perform various tasks to manage your deduplication disk pools.

See [“Viewing deduplication disk pools”](#) on page 162.

See [“Determining the deduplication disk pool state”](#) on page 162.

See [“Changing the deduplication disk pool state”](#) on page 162.

See [“Viewing deduplication disk pool attributes”](#) on page 163.

See [“Setting deduplication disk pool attributes”](#) on page 164.

See [“Changing deduplication disk pool properties”](#) on page 165.

See [“Clearing deduplication disk pool attributes”](#) on page 170.

See [“Determining the deduplication disk volume state”](#) on page 171.

See [“Changing the deduplication disk volume state”](#) on page 171.

See [“Deleting a deduplication disk pool”](#) on page 172.

Viewing deduplication disk pools

Use the NetBackup Administration Console to view configured disk pools.

To view disk pools

- ◆ In the NetBackup Administration Console, expand **Media and Device Management > Devices > Disk Pools**.

Determining the deduplication disk pool state

The disk pool state is UP or DOWN.

To determine disk pool state

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor**.
- 2 Select the **Disk Pools** tab.
- 3 The state is displayed in the **Status** column.

Changing the deduplication disk pool state

Disk pool state is UP or DOWN.

To change the state to DOWN, the disk pool must not be busy. If backup jobs are assigned to the disk pool, the state change fails. Cancel the backup jobs or wait until the jobs complete.

To change deduplication pool state

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 Select the **Disk Pools** tab.
- 3 Select the disk pool.
- 4 Select either **Actions > Up** or **Actions > Down**.

See [“About NetBackup deduplication pools”](#) on page 79.

See [“Determining the deduplication disk pool state”](#) on page 162.

See [“Media server deduplication pool properties”](#) on page 81.

See [“Configuring a deduplication disk pool”](#) on page 81.

Viewing deduplication disk pool attributes

Use the NetBackup `nbdevquery` command to view deduplication pool attributes.

To view deduplication pool attributes

- ◆ The following is the command syntax to view the attributes of a deduplication pool. Run the command on the NetBackup master server or on the deduplication storage server:

```
UNIX: /usr/opensv/netbackup/bin/admincmd/nbdevquery -listdp -dp  
pool_name -stype PureDisk -U
```

```
Windows: install_path\NetBackup\bin\admincmd\nbdevquery -listdp  
-dp pool_name -stype PureDisk -U
```

The following is example output:

```
Disk Pool Name      : MediaServerDeduplicationPool  
Disk Pool Id       : MediaServerDeduplicationPool  
Disk Type          : PureDisk  
Status             : UP  
Flag               : OpenStorage  
Flag               : AdminUp  
Flag               : InternalUp  
Flag               : LifeCycle  
Flag               : CapacityMgmt  
Flag               : OptimizedImage  
Raw Size (GB)     : 235.76  
Usable Size (GB)  : 235.76  
Num Volumes       : 1  
High Watermark    : 98  
Low Watermark     : 80  
Max IO Streams    : -1  
Storage Server    : DedupeServer.symantecs.org (UP)
```

This example output is shortened; more flags may appear in actual output.

Setting deduplication disk pool attributes

You may have to set attributes on your existing media server deduplication pools. For example, if you set an attribute on the storage server, you may have to set the same attribute on your existing deduplication disk pools.

See [“Setting deduplication storage server attributes”](#) on page 152.

To set a deduplication disk pool attribute

- 1 The following is the command syntax to set a deduplication pool attribute. Run the command on the master server or on the storage server.

```
nbdevconfig -changedp -dp pool_name -stype PureDisk -setattribute  
attribute
```

The following describes the options that require the arguments that are specific to your domain:

<code>-changedp</code> <i>pool_name</i>	The name of the disk pool.
<code>-setattribute</code> <i>attribute</i>	The <i>attribute</i> is the name of the argument that represents the new functionality. For example, OptimizedImage specifies that the environment supports the optimized synthetic backup method.

The following is the path to the `nbdevconfig` command:

- UNIX: `/usr/opensv/netbackup/bin/admincmd`
 - Windows: `install_path\NetBackup\bin\admincmd`
- 2 To verify, view the disk pool attributes.
See [“Viewing deduplication disk pool attributes”](#) on page 163.

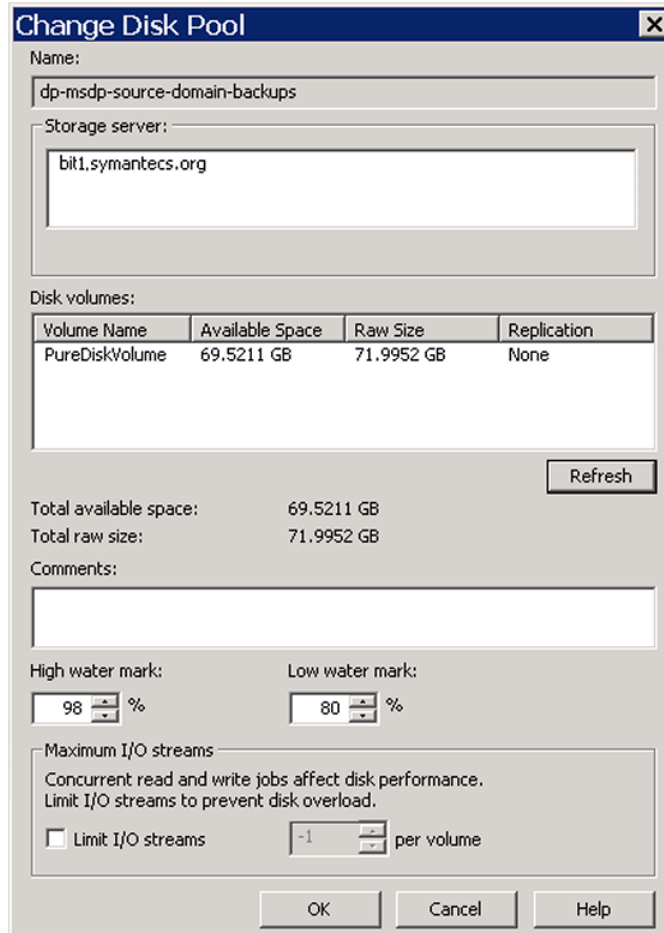
Changing deduplication disk pool properties

You can change the properties of a deduplication disk pool.

To change disk pool properties

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices > Disk Pools**.
- 2 Select the disk pool you want to change in the details pane.

- 3 On the **Edit** menu, select **Change**.



- 4 In the **Change Disk Pool** dialog box, click **Refresh** to update the disk pool replication properties.

If NetBackup discovers changes, your actions depend on the changes discovered.

See [“How to resolve volume changes for Auto Image Replication”](#) on page 167.

- 5 Change the other properties as necessary.

See [“Media server deduplication pool properties”](#) on page 81.

- 6 Click **OK**.
- 7 If you clicked **Refresh** and the **Replication** value for the **PureDiskVolume** changed, refresh the view in the **Administration Console**.

How to resolve volume changes for Auto Image Replication

When you open the **Change Disk Pool** dialog box, NetBackup loads the disk pool properties from the catalog. NetBackup only queries the storage server for changes when you click the **Refresh** in the **Change Disk Pool** dialog box.

Symantec recommends that you take the following actions when the volume topology change:

- Discuss the changes with the storage administrator. You need to understand the changes so you can change your disk pools (if required) so that NetBackup can continue to use them.
- If the changes were not planned for NetBackup, request that the changes be reverted so that NetBackup functions correctly again.

NetBackup can process changes to the following volume properties:

- Replication Source
- Replication Target
- None

If these volume properties change, NetBackup can update the disk pool to match the changes. NetBackup can continue to use the disk pool, although the disk pool may no longer match the storage unit or storage lifecycle purpose.

[Table 7-2](#) describes the possible outcomes and describes how to resolve them.

Table 7-2 Refresh outcomes

Outcome	Description
No changes are discovered.	No changes are required.
NetBackup discovers the new volumes that you can add to the disk pool.	The new volumes appear in the Change Disk Pool dialog box. Text in the dialog box changes to indicate that you can add the new volumes to the disk pool.

Table 7-2 Refresh outcomes (*continued*)

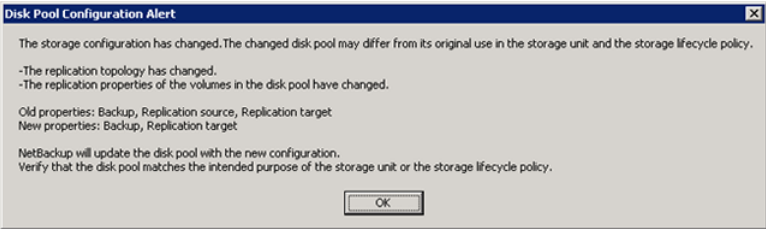
Outcome	Description
<p>The replication properties of all of the volumes changed, but they are still consistent.</p>	<p>A Disk Pool Configuration Alert pop-up box notifies you that the properties of all of the volumes in the disk pool changed, but they are all the same (homogeneous).</p>  <p>You must click OK in the alert box, after which the disk pool properties in the Change Disk Pool dialog box are updated to match the new volume properties.</p> <p>If new volumes are available that match the new properties, NetBackup displays those volumes in the Change Disk Pool dialog box. You can add those new volumes to the disk pool.</p> <p>In the Change Disk Pool dialog box, select one of the following two choices:</p> <ul style="list-style-type: none"> ■ OK. To accept the disk pool changes, click OK in the Change Disk Pool dialog box. NetBackup saves the new properties of the disk pool. NetBackup can use the disk pool, but it may no longer match the intended purpose of the storage unit or storage lifecycle policy. Change the storage lifecycle policy definitions to ensure that the replication operations use the correct source and target disk pools, storage units, and storage unit groups. Alternatively, work with your storage administrator to change the volume properties back to their original values. ■ Cancel. To discard the changes, click Cancel in the Change Disk Pool dialog box. NetBackup does not save the new disk pool properties. NetBackup can use the disk pool, but it may no longer match the intended use of the storage unit or storage lifecycle policy.

Table 7-2 Refresh outcomes (continued)

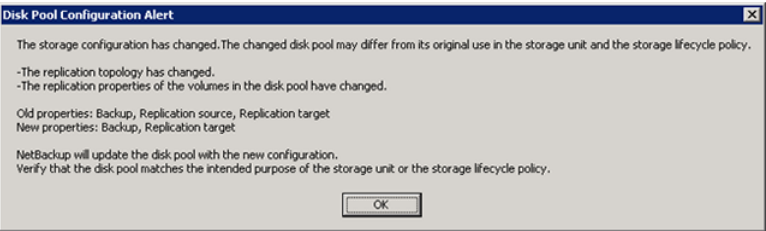
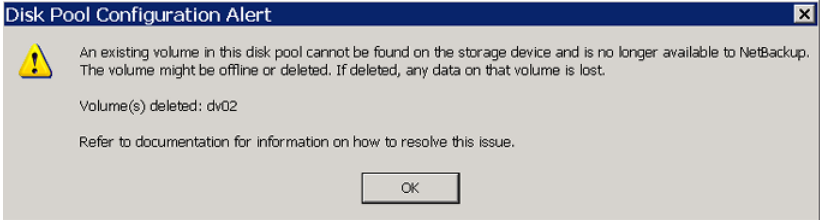
Outcome	Description
<p>The replication properties of the volumes changed, and they are now inconsistent.</p>	<p>A Disk Pool Configuration Error pop-up box notifies you that the replication properties of some of the volumes in the disk pool changed. The properties of the volumes in the disk pool are not homogeneous.</p>  <p>You must click OK in the alert box.</p> <p>In the Change Disk Pool dialog box, the properties of the disk pool are unchanged, and you cannot select them (that is, they are dimmed). However, the properties of the individual volumes are updated.</p> <p>Because the volume properties are not homogeneous, NetBackup cannot use the disk pool until the storage configuration is fixed.</p> <p>NetBackup does not display new volumes (if available) because the volumes already in the disk pool are not homogeneous.</p> <p>To determine what has changed, compare the disk pool properties to the volume properties.</p> <p>See “Viewing the replication topology for Auto Image Replication” on page 100.</p> <p>Work with your storage administrator to change the volume properties back to their original values.</p> <p>The disk pool remains unusable until the properties of the volumes in the disk pool are homogenous.</p> <p>In the Change Disk Pool dialog box, click OK or Cancel to exit the Change Disk Pool dialog box.</p>

Table 7-2 Refresh outcomes (*continued*)

Outcome	Description
NetBackup cannot find a volume or volumes that were in the disk pool.	<p>A Disk Pool Configuration Alert pop-up box notifies you that an existing volume or volumes was deleted from the storage device:</p>  <p>NetBackup can use the disk pool, but data may be lost.</p> <p>To protect against accidental data loss, NetBackup does not allow volumes to be deleted from a disk pool.</p> <p>To continue to use the disk pool, do the following:</p> <ul style="list-style-type: none"> ■ Use the <code>bpimmedia</code> command or the Images on Disk report to display the images on the specific volume. ■ Expire the images on the volume. ■ Use the <code>nbdevconfig</code> command to set the volume state to DOWN so NetBackup does not try to use it.

Clearing deduplication disk pool attributes

You may have to clear attributes on your existing media server deduplication pools.

To clear a Media Server Deduplication Pool attribute

- ◆ The following is the command syntax to clear a deduplication pool attribute. Run the command on the master server or on the storage server.

```
nbdevconfig -changedp -dp pool_name -stype PureDisk
-clearattribute attribute
```

The following describe the options that require your input:

<code>-changedp</code> <i>pool_name</i>	The name of the disk pool.
<code>-setattribute</code> <i>attribute</i>	The <i>attribute</i> is the name of the argument that represents the new functionality.

The following is the path to the `nbdevconfig` command:

- **UNIX:** `/usr/opensv/netbackup/bin/admincmd`
- **Windows:** `install_path\NetBackup\bin\admincmd`

Determining the deduplication disk volume state

Use the `NetBackup nbdevquery` command to determine the state of the volume in a deduplication disk pool. The command shows the properties and attributes of the `PureDiskVolume`.

To determine deduplication disk volume state

- ◆ Display the volume state by using the following command:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbdevquery -listdv -stype PureDisk -U`

Windows: `install_path\NetBackup\bin\admincmd\nbdevquery -listdv -stype PureDisk -U`

The *state* is either UP or DOWN.

The following is example output

```

Disk Pool Name       : PD_Disk_Pool
Disk Type            : PureDisk
Disk Volume Name     : PureDiskVolume
Disk Media ID        : @aaaab
Total Capacity (GB) : 49.98
Free Space (GB)      : 43.66
Use%                  : 12
Status                : UP
Flag                  : ReadOnWrite
Flag                  : AdminUp
Flag                  : InternalUp
Num Read Mounts      : 0
Num Write Mounts     : 1
Cur Read Streams    : 0
Cur Write Streams   : 0
    
```

Changing the deduplication disk volume state

The disk volume state is UP or DOWN.

To change the state to DOWN, the disk pool in which the volume resides must not be busy. If backup jobs are assigned to the disk pool, the state change fails. Cancel the backup jobs or wait until the jobs complete.

To change the deduplication disk volume state

- 1 Determine the name of the disk volume. The following command lists all volumes in the specified disk pool:

```
nbdevquery -listdv -stype PureDisk -dp disk_pool_name
```

The `nbdevquery` and the `nbdevconfig` commands reside in the following directory:

- UNIX: `/usr/opensv/NetBackup/bin/admincmd`
- Windows: `install_path\NetBackup\bin\admincmd`

To display the disk volumes in all disk pools, omit the `-dp` option.

- 2 Change the disk volume state; the following is the command syntax:

```
nbdevconfig -changestate -stype PureDisk -dp disk_pool_name -dv  
vol_name -state state
```

The *state* is either UP or DOWN.

Deleting a deduplication disk pool

You can delete a disk pool if it does not contain valid NetBackup backup images or image fragments. If it does, you must first expire and delete those images or fragments. If expired image fragments remain on disk, you must remove those also.

See “[Expired fragments remain on disk](#)” on page 198.

If you delete a disk pool, NetBackup removes it from your configuration.

If a disk pool is the storage destination of a storage unit, you must first delete the storage unit.

To delete a deduplication disk pool

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Disk Pools**.
- 2 Select a disk pool
- 3 On the **Edit** menu, select **Delete**.
- 4 In the **Delete Disk Pool** dialog box, verify that the disk pool is the one you want to delete and then click **OK**.

Deleting backup images

Image deletion may be time consuming. Therefore, if you delete images manually, Symantec recommends the following approach.

See “[Data removal process](#)” on page 224.

To delete backup images manually

- 1 Expire all of the images by using the `bpexpdate` command and the `-notimmediate` option. The `-notimmediate` option prevents `bpexpdate` from calling the `nbdelete` command, which deletes the image.

Without this option, `bpexpdate` calls `nbdelete` to delete images. Each call to `nbdelete` creates a job in the Activity Monitor, allocates resources, and launches processes on the media server.

- 2 After you expire the last image, delete all of the images by using the `nbdelete` command with the `-allvolumes` option.

Only one job is created in the Activity Monitor, fewer resources are allocated, and fewer processes are started on the media servers. The entire process of expiring images and deleting images takes less time.

Disabling client-side deduplication for a client

You can remove a client from the list of clients that deduplicate their own data. If you do so, a deduplication server backs up the client and deduplicates the data.

To disable client deduplication for a client

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the details pane, select the master server.
- 3 On the **Actions** menu, select **Properties**.
- 4 On the **Host Properties Client Attributes General** tab, select the client that deduplicates its own data.
- 5 In the **Deduplication Location** drop-down list, select **Always use the media server**.
- 6 Click **OK**.

About deduplication queue processing

Operations that require database updates accumulate in a transaction queue. Twice a day, the NetBackup Deduplication Manager directs the deduplication engine to process the queue as one batch.

The schedule is frequency-based. By default, queue processing occurs every 12 hours, 20 minutes past the hour.

Queue processing consumes two CPU cores.

Queue processing writes status information to the deduplication engine `storaged.log` file.

See [“About deduplication logs”](#) on page 187.

In a few rare scenarios, some data segments may become orphaned. Queue processing includes operations that remove those segments. (Previously, this was a separate operation known as garbage collection.)

Because queue processing does not block any other deduplication process, rescheduling should not be necessary. Users cannot change the maintenance process schedules. However, if you must reschedule these processes, contact your Symantec support representative.

Because queue processing occurs automatically, you should not need to invoke it manually. However, you may do so.

See [“Processing the deduplication transaction queue manually”](#) on page 174.

See [“About deduplication server requirements”](#) on page 26.

Processing the deduplication transaction queue manually

Usually, you should not need to run the deduplication database transaction queue processes manually. However, you can do so.

See [“About deduplication queue processing”](#) on page 174.

A control command launches the queue processing. The following is the path name of the command:

- On UNIX and Linux systems, `/usr/opensv/pdde/pdcr/bin/crcontrol`.
- On Windows systems, `install_path\Veritas\pdde\Crcontrol.exe`.

To perform deduplication maintenance manually

- 1 Run the control command with the `--processqueue` option. The following is an example on a Windows system:

```
install_path\Veritas\pdde\Crcontrol.exe --processqueue
```

- 2 To examine the results, run the control command with the `--dsstat 1` option (number 1 not lowercase letter l). The command may run for a long time; if you omit the 1, results return more quickly but they are not as accurate.

See [“Viewing storage usage within deduplication container files”](#) on page 143.

About deduplication data integrity checking

Deduplication metadata and data may become inconsistent or corrupted because of disk failures, I/O errors, database corruption, and operational errors. NetBackup checks the integrity of the deduplicated data on a regular basis. NetBackup performs some of the integrity checking when the storage server is idle. Other integrity checking is designed to use few storage server resources so as not to interfere with operations.

The data integrity checking process includes the following checks:

- Data consistency check
- Cyclic redundancy check (CRC)
- Storage leak check

NetBackup resolves many integrity issues without user intervention, and some issues are fixed when the next backup runs. However, a severe issue may require intervention by Symantec Support. In such cases, NetBackup writes a message to the NetBackup Disk Logs report.

See [“Viewing disk reports”](#) on page 145.

Data integrity message codes are 1047, 1057, and 1058.

See [“Deduplication event codes and messages”](#) on page 202.

NetBackup writes integrity checking activity messages to the `storaged.log` file.

See [“About deduplication logs”](#) on page 187.

You can configure some of the data integrity checking behaviors.

See [“Configuring deduplication data integrity checking behavior”](#) on page 176.

Configuring deduplication data integrity checking behavior

NetBackup performs three different data integrity checks. You can configure the behavior of those checks.

Two methods exist to configure deduplication data integrity checking behavior, as follows:

- Edit a configuration file parameter.
See [“To configure data integrity checking behavior by editing the configuration file”](#) on page 176.
- Run a command.
See [“To configure data integrity checking behavior by using a command”](#) on page 176.

More information about data integrity checking is available.

See [“About deduplication data integrity checking”](#) on page 175.

See [“Deduplication data integrity checking configuration settings”](#) on page 178.

To configure data integrity checking behavior by editing the configuration file

- 1 Use a text editor to open the `contentrouter.cfg` file.

The `contentrouter.cfg` file resides in the following directories:

- UNIX: `storage_path/etc/puredisk`
- Windows: `storage_path\etc\puredisk`

- 2 To change a parameter, specify a new value.

Note: The spaces to the left and right of the equal sign (=) in the file are significant. Ensure that the space characters appear in the file after you edit the file.

See [“Deduplication data integrity checking configuration settings”](#) on page 178.

- 3 Save and close the file.

To configure data integrity checking behavior by using a command

- ◆ To configure behavior, specify a value for each of the data integrity checks, as follows:
 - Data consistency checking. Use the following commands to configure behavior:

Configuring deduplication data integrity checking behavior

Enable	UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a enabledataintegritycheck</code>
Disable	UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a disabledataintegritycheck</code>
Get the status	UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a getdataintegritycheck</code>

■ **CRC checking. Use the following commands to configure behavior:**

Enable	<p>CRC check does not run if queue processing is active or during disk read or write operations.</p> <p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckon</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --crccheckon</code></p>
Disable	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckoff</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --crccheckoff</code></p>
Get the status	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckstate</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --crccheckstate</code></p>

■ **Storage leak checking. Use the following commands to configure behavior:**

Enable	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --storageleakcheckon</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --storageleakcheckon</code></p>
--------	--

Disable	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --storageleakcheckoff</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --storageleakcheckoff</code></p>
Get the status	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --storageleakcheckstate</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --storageleakcheckstate</code></p>

Deduplication data integrity checking configuration settings

The following are the configuration file parameters that you can set for deduplication data integrity checking. The parameters are in the `contentrouter.cfg` file and the `spa.cfg` file. Those files reside in the following directories:

- UNIX: `storage_path/etc/puredisk`
- Windows: `storage_path\etc\puredisk`

Table 7-3 spa.cfg file parameters for data integrity checking

Setting	Default value	Description
EnableDataCheck	True	Enable or disable data consistency checking. The possible values are True or False.
DataCheckDays	30	The number of days to check the data for consistency. The greater the number of days, the fewer the objects that are checked each day. The greater the number of days equals fewer storage server resources consumed each day.

Table 7-4 contentrouter.cfg file parameters for data integrity checking

Setting	Default value	Description
EnableCRCCheck	True	Enable or disable CRC checking of the data container files. The possible values are <code>True</code> or <code>False</code> . CRC checking occurs only when no backup, restore, or queue processing jobs are running.
CRCCheckSleepSeconds	2	The time in seconds to sleep between checking containers. The longer the sleep interval, the more time it takes to check containers.
EnableStorageLeakCheck	True	Enable or disable storage leak checking. The possible values are <code>True</code> or <code>False</code> .
CheckExpirationDays	60	The number of days to test for storage leaks. Usually, set this value to double that of the <code>DataCheckDays</code> value. The greater the number of days, the fewer the objects that are checked each day. The greater the number of days equals fewer storage server resources consumed each day.

About managing storage read performance

NetBackup provides some control over the processes that are used for read operations. The read operation controls can improve performance for the jobs that read from the storage. Such jobs include restore jobs and duplication jobs to tape. Duplication jobs to tape rehydrate the data into the original NetBackup backup images and write the images to tape for long-term archiving.

In most cases, you should change configuration file options only when directed to do so by Symantec Technical Support.

Table 7-5 Some settings that affect restore operations

Process	Description
Defragmenting the storage	<p>NetBackup includes a process, called <i>rebasing</i>, which defragments the backup images in a deduplication pool. Read performance improves when the file segments from a client backup are close to each other on deduplication storage.</p> <p>If you upgrade from a NetBackup release earlier than 7.5, rebasing may affect your deduplication performance temporarily.</p> <p>See “About deduplication storage rebasing” on page 180.</p>
The number of prefetching threads	<p>The <code>PrefetchThreadNum</code> parameter in the <code>contentrouter.cfg</code> file specifies the number of threads to use to preload segments during storage read operations.</p> <p>The default value is 1. Depending on your disks, a value as high as 4 may improve performance. However, Symantec recommends that you test higher values thoroughly to ensure that a value greater than 1 yields better performance. Higher values may actually decrease performance.</p> <p>See “About the contentrouter.cfg file for NetBackup deduplication” on page 127.</p>
The prefetch buffer size	<p>The <code>PREFETCH_SIZE</code> parameter in the <code>pd.conf</code> file specifies the size in bytes to use for the data buffer for read operations.</p> <p>See “About the pd.conf configuration file for NetBackup deduplication” on page 119.</p> <p>See “pd.conf file settings for NetBackup deduplication” on page 119.</p>
Decrypting the data on the client rather than the server	<p>The <code>RESTORE_DECRYPT_LOCAL</code> parameter in the <code>pd.conf</code> file specifies on which host to decrypt and decompress the data during restore operations.</p> <p>See “About the pd.conf configuration file for NetBackup deduplication” on page 119.</p> <p>See “pd.conf file settings for NetBackup deduplication” on page 119.</p>

About deduplication storage rebasing

NetBackup controls data segment locality during backups, compaction, and optimized duplication operations. However, a file's segments may become scattered across the disk storage as the file continues to change and continues to be backed up. Such scattering is a normal consequence of deduplication.

NetBackup includes a process, called *rebasing*, which defragments the backup images in a deduplication pool. Read performance improves when the file segments from a client backup are close to each other on deduplication storage. NetBackup consumes less time finding and reassembling files when their segments are near each other.

Table 7-6 Types of rebasing

Rebasing type	Description
Server-side rebasing	<p>NetBackup automatically rebases deduplicated data in a NetBackup Media Server Deduplication Pool. The rebasing process is configured to occur when the storage server is not otherwise busy.</p> <p>The following parameters in the <code>contentrouter.cfg</code> file control server-side rebasing behavior:</p> <ul style="list-style-type: none"> ■ RebaseScatterThreshold This parameter specifies the average data size threshold per container for a given backup image to be considered for rebasing. By default, this parameter is <code>RebaseScatterThreshold=64MB</code>. ■ RebaseQuota This parameter specifies the amount of data that the rebasing operation can relocate (or move) each day. By default, this parameter is <code>RebaseQuota=500GB</code>. <p>Usually, rebasing has little affect on normal operations. However, immediately after an upgrade to NetBackup 7.5, rebasing may overlap into your backup windows. If performance degrades unacceptably during rebasing, change the parameter values in the <code>contentrouter.cfg</code> file. A lower value for each parameter reduces the resources that the storage server uses for rebasing.</p> <p>See “About the contentrouter.cfg file for NetBackup deduplication” on page 127.</p>

Table 7-6 Types of rebasing (*continued*)

Rebasing type	Description
Client-side rebasing	<p>Client-side deduplication only.</p> <p>Client-side rebasing can yield further performance increases. Segments from a backup of the specified client with the specified policy are sent to the storage server without checking for the existence of a segment on the specified date. In the deduplication pool, these segments are stored next to each other within containers. Queue processing removes previously stored, duplicated segments to keep unique segments within the container.</p> <p>Some limitations exist, as follows:</p> <ul style="list-style-type: none"> ■ The deduplication rate per job may not accurately reflect the actual storage deduplication rate. ■ Performance may not improve for the remote clients that have limited network bandwidth. <p>The <code>CLIENT_POLICY_DATE</code> parameter in the <code>pd.conf</code> file controls rebasing on a NetBackup client.</p> <p>See “About the pd.conf configuration file for NetBackup deduplication” on page 119.</p> <p>See “pd.conf file settings for NetBackup deduplication” on page 119.</p> <p>NetBackup automatically adds the <code>CLIENT_POLICY_DATE</code> parameter to Linux and UNIX NetBackup client-side deduplication clients. If rebasing is required on Windows clients, edit the <code>pd.conf</code> file and add the parameter manually.</p>

Resizing the deduplication storage partition

If the volume that contains the deduplication storage is resized dynamically, restart the NetBackup services on the storage server. You must restart the services so that NetBackup can use the resized partition correctly. If you do not restart the services, NetBackup reports the capacity as full prematurely.

To resize the deduplication storage

- 1 Stop all NetBackup jobs on the storage on which you want to change the disk partition sizes and wait for the jobs to end.
- 2 Deactivate the media server that hosts the storage server.
See the *NetBackup Administrator's Guide for UNIX and Linux, Volume I* or the *NetBackup Administrator's Guide for Windows, Volume I*.
- 3 Stop the NetBackup services on the storage server.
Be sure to wait for all services to stop.
- 4 Use the operating system or disk manager tools to dynamically increase or decrease the deduplication storage area.

- 5 Restart the NetBackup services.
- 6 Activate the media server that hosts the storage server.
See the *NetBackup Administrator's Guide for UNIX and Linux, Volume I* or the *NetBackup Administrator's Guide for Windows, Volume I*.
- 7 Restart the deduplication jobs.

About restoring files at a remote site

If you use optimized duplication to copy images from a local site to a remote site, you can restore from the copies at the remote site to clients at the remote site. To do so, use a server-directed restore or a client-redirected restore, which restores files to a client other than the original client.

Information about how to redirect restores is in a different guide.

See “Managing client restores” in the *NetBackup Administrator's Guide for UNIX and Linux, Volume I* or the *NetBackup Administrator's Guide for Windows, Volume I*.

You may have to configure which media server performs the restore. In optimized duplication, the media server that initiates the duplication operation becomes the write host for the new image copies. The write host restores from those image copies. If the write host is at the local site, it restores from those images at the remote site to the alternate client at the remote site. That host reads the image across the WAN and then writes the image back across the WAN to the alternate client. In this case, you can specify that the media server at the remote site as the restore server.

About restoring from a backup at a target master domain

While it is possible to restore a client directly by using the images in the target master domain, do so only in a disaster recovery situation. In this discussion, a disaster recovery situation is one in which the originating domain no longer exists and clients must be recovered from the target domain.

Table 7-7 Client restores in disaster recovery scenarios

Disaster recovery scenario	Does client exist?	Description
Scenario 1	Yes	Configure the client in another domain and restore directly to the client.
Scenario 2	No	Create the client in the recovery domain and restore directly to the client. This is the most likely scenario.
Scenario 3	No	Perform an alternate client restore in the recovery domain.

The steps to recover the client are the same as any other client recovery. The actual steps depend on the client type, the storage type, and whether the recovery is an alternate client restore.

For restores that use Granular Recovery Technology (GRT), an application instance must exist in the recovery domain. The application instance is required so that NetBackup has something to recover to.

Specifying the restore server

NetBackup may not use the backup server as the restore server for deduplicated data.

See [“How deduplication restores work”](#) on page 59.

You can specify the server to use for restores. The following are the methods that specify the restore server:

- Always use the backup server. Two methods exist, as follows:
 - Use NetBackup **Host Properties** to specify a **Media host override** server. All restore jobs for any storage unit on the original backup server use the media server you specify. Specify the same server for the **Restore server** as for the **Original backup server**.
 See “Forcing restores to use a specific server” in the *NetBackup Administrator's Guide for UNIX and Linux, Volume I* or the *NetBackup Administrator's Guide for Windows, Volume I*.
 This procedure sets the `FORCE_RESTORE_MEDIA_SERVER` option. Configuration options are stored in the `bp.conf` file on UNIX systems and the registry on Windows systems.
 - Create the touch file `USE_BACKUP_MEDIA_SERVER_FOR_RESTORE` on the NetBackup master server in the following directory:

UNIX: `usr/opensv/netbackup/db/config`

Windows: `install_path\veritas\netbackup\db\config`

This global setting always forces restores to the server that did the backup. It applies to all NetBackup restore jobs, not just deduplication restore jobs. If this touch file exists, NetBackup ignores the

`FORCE_RESTORE_MEDIA_SERVER` and `FAILOVER_RESTORE_MEDIA_SERVER` settings.

- Always use a different server.
Use NetBackup **Host Properties** to specify a **Media host override** server. See the previous explanation about **Media host override**, except: Specify the different server for the **Restore server**.
- A single restore instance. Use the `bprestore` command with the `-disk_media_server` option.
Restore jobs for each instance of the command use the media server you specify. See *NetBackup Commands Reference Guide*.

Troubleshooting

This chapter includes the following topics:

- [About deduplication logs](#)
- [Troubleshooting installation issues](#)
- [Troubleshooting configuration issues](#)
- [Troubleshooting operational issues](#)
- [Viewing disk errors and events](#)
- [Deduplication event codes and messages](#)

About deduplication logs

The NetBackup deduplication components write information to various log files.

[Table 8-1](#) describes the log files for each component.

Information about VxUL log files is available.

See “[About VxUL logs for deduplication](#)” on page 190.

Table 8-1 Logs for NetBackup deduplication activity

Component	Description
Client deduplication proxy plug-in	<p>The client deduplication proxy plug-in on the media server runs under <code>bptm</code>, <code>bpstsinfo</code>, and <code>bpbrm</code> processes. Examine the log files for those processes for proxy plug-in activity. The strings <code>proxy</code> or <code>ProxyServer</code> embedded in the log messages identify proxy server activity.</p> <p>They write log files to the following directories:</p> <ul style="list-style-type: none"> ■ For <code>bptm</code>: <ul style="list-style-type: none"> UNIX: <code>/usr/opensv/netbackup/logs/bptm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bptm</code> ■ For <code>bpstsinfo</code>: <ul style="list-style-type: none"> Windows: <code>/usr/opensv/netbackup/logs/admin</code> UNIX: <code>/usr/opensv/netbackup/logs/bpstsinfo</code> Windows: <code>install_path\Veritas\NetBackup\logs\admin</code> Windows: <code>install_path\Veritas\NetBackup\logs\stsinfo</code> ■ For <code>bpbrm</code>: <ul style="list-style-type: none"> UNIX: <code>/usr/opensv/netbackup/logs/bpbrm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bpbrm</code>
Client deduplication proxy server	<p>The deduplication proxy server <code>nbostpxy</code> on the client writes messages to files in an eponymous directory, as follows:</p> <p>UNIX: <code>/usr/opensv/netbackup/logs/nbostpxy</code></p> <p>Windows: <code>install_path\Veritas\NetBackup\logs\nbostpxy</code>.</p>
Deduplication configuration script	<p>The following is the path name of the log file for the deduplication configuration script:</p> <ul style="list-style-type: none"> ■ UNIX: <code>storage_path/log/pdde-config.log</code> ■ Windows: <code>storage_path\log\pdde-config.log</code> <p>NetBackup creates this log file during the configuration process. If your configuration succeeded, you do not need to examine the log file. The only reason to look at the log file is if the configuration failed. If the configuration process fails after it creates and populates the storage directory, this log file identifies when the configuration failed.</p>
Deduplication database	<p>The deduplication database log file (<code>postgresql.log</code>) is in the <code>storage_path/databases/pddb</code> directory.</p> <p>You can configure log parameters. For more information, see the following: http://www.postgresql.org/docs/current/static/runtime-config-logging.html</p> <p>See “Configuring deduplication log file timestamps on Windows” on page 133.</p>

Table 8-1 Logs for NetBackup deduplication activity (*continued*)

Component	Description
NetBackup Deduplication Engine	<p>The NetBackup Deduplication Engine writes several log files, as follows:</p> <ul style="list-style-type: none"> ■ Log files in the <code>storage_path/log/spoold</code> directory, as follows: <ul style="list-style-type: none"> ■ The <code>spoold.log</code> file is the main log file ■ The <code>stored.log</code> file is for queue processing. ■ A log file for each connection to the engine is stored in a directory in the storage path <code>spoold</code> directory. The following describes the pathname to a log file for a connection: <code>hostname/application/TaskName/MMDDYY.log</code> For example, the following is an example of a <code>crcontrol</code> connection log pathname on a Linux system: <code>/storage_path/log/spoold/server.symantecs.org/crcontrol/Control/010112.log</code> Usually, the only reason to examine these connection log files is if a Symantec support representative asks you to. ■ A VxUL log file for the events and errors that NetBackup receives from polling. The originator ID for the deduplication engine is 364. See “About VxUL logs for deduplication” on page 190.
NetBackup Deduplication Manager	<p>The log files are in the <code>/storage_path/log/spad</code> directory, as follows:</p> <ul style="list-style-type: none"> ■ <code>spad.log</code> ■ <code>sched_QueueProcess.log</code> ■ <code>SchedClass.log</code> ■ A log file for each connection to the manager is stored in a directory in the storage path <code>spad</code> directory. The following describes the pathname to a log file for a connection: <code>hostname/application/TaskName/MMDDYY.log</code> For example, the following is an example of a <code>bpstsinfo</code> connection log pathname on a Linux system: <code>/storage_path/log/spoold/server.symantecs.org/bpstsinfo/spad/010112.log</code> Usually, the only reason to examine these connection log files is if a Symantec support representative asks you to. <p>You can set the log level and retention period in the Change Storage Server dialog box Properties tab.</p> <p>See “Changing deduplication storage server properties” on page 153.</p>
Deduplication plug-in	<p>You can configure the location and name of the log file and the logging level. To do so, edit the <code>DEBUGLOG</code> entry and the <code>LOGLEVEL</code> in the <code>pd.conf</code> file.</p> <p>See “About the pd.conf configuration file for NetBackup deduplication” on page 119.</p> <p>See “Editing the pd.conf deduplication file” on page 126.</p>

About VxUL logs for deduplication

Some NetBackup commands or processes write messages to their own log files. Other processes use Veritas unified log (VxUL) files. VxUL uses a standardized name and file format for log files. An originator ID (OID) identifies the process that writes the log messages.

[Table 8-2](#) shows the NetBackup logs for disk-related activity.

The messages that begin with a `sts_` prefix relate to the interaction with the storage vendor software plug-in. Most interaction occurs on the NetBackup media servers.

To view and manage VxUL log files, you must use NetBackup log commands. For information about how to use and manage logs on NetBackup servers, see the *NetBackup Troubleshooting Guide*.

Table 8-2 NetBackup VxUL logs

Activity	VxUL OID	Processes that use the ID
NetBackup Deduplication Engine	364	The NetBackup Deduplication Engine that runs on the deduplication storage server.
Backups and restores	N/A	Messages appear in the log files for the following processes: <ul style="list-style-type: none"> ■ The <code>bpbm</code> backup and restore manager ■ The <code>bpbm</code> database manager ■ The <code>bptm</code> tape manager for I/O operations
Backups and restores	117	The <code>nbjm</code> Job Manager.
Device configuration and monitoring	111	The <code>nbeem</code> process.
Device configuration and monitoring	178	The Disk Service Manager process that runs in the Enterprise Media Manager (EMM) process.
Device configuration and monitoring	202	The storage server interface process that runs in the Remote Manager and Monitor Service. RMMS runs on media servers.
Device configuration and monitoring	230	The Remote Disk Service Manager interface (RDSM) that runs in the Remote Manager and Monitor Service. RMMS runs on media servers.

Table 8-2 NetBackup VxUL logs (*continued*)

Activity	VxUL OID	Processes that use the ID
Resilient network connections	387	The Remote Network Transport Service (<code>nbrntd</code>) manages resilient network connections. It runs on the master server, on media servers, and on clients.

Troubleshooting installation issues

The following sections may help you troubleshoot installation issues.

See [“Installation on SUSE Linux fails”](#) on page 191.

Installation on SUSE Linux fails

The installation trace log shows an error when you install on SUSE Linux:

```
...NetBackup and Media Manager are normally installed in /usr/opensv.
Is it OK to install in /usr/opensv? [y,n] (y)

Reading NetBackup files from /net/nbstore/vol/test_data/PDDE_packages/
suse/NB_FID2740_LinuxS_x86_20090713_6.6.0.27209/linuxS_x86/anb

/net/nbstore/vol/test_data/PDDE_packages/suse/NB_FID2740_LinuxS_x86_
20090713_6.6.0.27209/linuxS_x86/catalog/anb/NB.file_trans: symbol
lookup error: /net/nbstore/vol/test_data/PDDE_packages/suse/
NB_FID2740_LinuxS_x86_20090713_6.6.0.27209/linuxS_x86/catalog/anb/
NB.file_trans: undefined symbol: head /net/nbstore/vol/test_data/
PDDE_packages/suse/NB_FID2740_LinuxS_x86_20090713_6.6.0.27209/
linuxS_x86/catalog/anb/NB.file_trans failed. Aborting ...
```

Verify that your system is at patch level 2 or later, as follows:

```
cat /etc/SuSE-release
SUSE Linux Enterprise Server 10 (x86_64)
VERSION = 10
PATCHLEVEL = 2
```

Troubleshooting configuration issues

The following sections may help you troubleshoot configuration issues.

See [“About deduplication logs”](#) on page 187.

See [“Storage server configuration fails”](#) on page 192.

See [“Database system error \(220\)”](#) on page 192.

See [“Server not found error”](#) on page 193.

See [“License information failure during configuration”](#) on page 193.

See [“The disk pool wizard does not display a volume”](#) on page 194.

Storage server configuration fails

If storage server configuration fails, first resolve the issue that the **Storage Server Configuration Wizard** reports. Then, delete the deduplication host configuration file before you try to configure the storage server again.

NetBackup cannot configure a storage server on a host on which a storage server already exists. One indicator of a configured storage server is the deduplication host configuration file. Therefore, it must be deleted before you try to configure a storage server after a failed attempt.

See [“Deleting a deduplication host configuration file”](#) on page 132.

Database system error (220)

A database system error indicates that an error occurred in the storage initialization.

Error message `ioctl() error, Database system error (220)`

Example RDSM has encountered an STS error:

```
Failed to update storage server ssname, database
system error
```

Diagnosis The `PDDE_initConfig` script was invoked, but errors occurred during the storage initialization.

First, examine the deduplication configuration script log file for references to the server name.

See [“About deduplication logs”](#) on page 187..

Second, examine the `tpconfig` command log file errors about creating the credentials for the server name. The `tpconfig` command writes to the standard NetBackup administrative commands log directory.

Server not found error

The following information may help you resolve a server not found error message that may occur during configuration.

Error message	Server not found, invalid command parameter
Example	<pre>RDSM has encountered an issue with STS where the server was not found: getStorageServerInfo Failed to create storage server <i>ssname</i>, invalid command parameter</pre>
Diagnosis	<p>Possible root causes:</p> <ul style="list-style-type: none"> ■ When you configured the storage server, you selected a media server that runs an unsupported operating system. All media servers in your environment appear in the Storage Server Configuration Wizard; be sure to select only a media server that runs a supported operating system. ■ If you used the <code>nbdevconfig</code> command to configure the storage server, you may have typed the host name incorrectly. Also, case matters for the storage server type, so ensure that you use PureDisk for the storage server type.

License information failure during configuration

A configuration error message about license information failure indicates that the NetBackup servers cannot communicate with each other.

If you cannot configure a deduplication storage server or load balancing servers, your network environment may not be configured for DNS reverse name lookup.

You can edit the hosts file on the media servers that you use for deduplication. Alternatively, you can configure NetBackup so it does not use reverse name lookup.

To prohibit reverse host name lookup by using the Administration Console

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the details pane, select the master server.
- 3 On the **Actions** menu, select **Properties**.
- 4 In the **Master Server Properties** dialog box, select the **Network Settings** properties.
- 5 Select one of the following options:

- **Allowed**
- **Restricted**
- **Prohibited**

For a description of these options, see the NetBackup online Help or the administrator's guide.

To prohibit reverse host name lookup by using the `bpsetconfig` command

- ◆ Enter the following command on each media server that you use for deduplication:

```
echo REVERSE_NAME_LOOKUP = PROHIBITED | bpsetconfig -h host_name
```

The `bpsetconfig` command resides in the following directories:

UNIX: `/usr/opensv/netbackup/bin/admincmd`

Windows: `install_path\Veritas\NetBackup\bin\admincmd`

The disk pool wizard does not display a volume

The **Disk Pool Configuration Wizard** does not display a disk volume for the deduplication storage server.

First, restart all of the NetBackup daemons or services. The step ensures that the NetBackup Deduplication Engine is up and ready to respond to requests.

Second, restart the **NetBackup Administration Console**. This step clears cached information from the failed attempt to display the disk volume.

Troubleshooting operational issues

The following sections may help you troubleshoot operational issues.

See [“Verify that the server has sufficient memory”](#) on page 195.

See [“Backup or duplication jobs fail”](#) on page 195.

See [“Client deduplication fails”](#) on page 196.

See [“Volume state changes to DOWN when volume is unmounted”](#) on page 197.

See [“Errors, delayed response, hangs”](#) on page 198.

See [“Cannot delete a disk pool”](#) on page 198.

See [“Media open error \(83\)”](#) on page 199.

See [“Media write error \(84\)”](#) on page 200.

See [“Storage full conditions”](#) on page 201.

See [“Specifying the restore server”](#) on page 184.

Verify that the server has sufficient memory

Insufficient memory on the storage server can cause operation problems. If you have operation issues, you should verify that your storage server has sufficient memory.

See [“About deduplication server requirements”](#) on page 26.

If the NetBackup deduplication processes do not start on Red Hat Linux, configure shared memory to be at least 128 MB (`SHMMAX=128MB`).

Backup or duplication jobs fail

[Table 8-3](#) describes some potential failures for backup or deduplication jobs and how to resolve them.

Table 8-3 Backup or deduplication job failures

Error condition	Description
Error 800: Disk Volume is Down	<p>Examine the disk error logs to determine why the volume was marked DOWN.</p> <p>If the storage server is busy with jobs, it may not respond to master server disk polling requests in a timely manner. A busy load balancing server also may cause this error. Consequently, the query times out and the master server marks the volume DOWN.</p> <p>If the error occurs for an optimized duplication job: verify that source storage server is configured as a load balancing server for the target storage server. Also verify that the target storage server is configured as a load balancing server for the source storage server.</p> <p>See “Viewing disk errors and events” on page 202.</p>

Table 8-3 Backup or deduplication job failures (*continued*)

Error condition	Description
<p>Error nbjm(pid=6384) NBU status: 2106, EMM status: Storage Server is down or unavailable Disk storage server is down(2106)</p>	<p>Windows servers only.</p> <p>The NetBackup Deduplication Manager (<i>spad.exe</i>) and the NetBackup Deduplication Engine (<i>spoold.exe</i>) have different shared memory configuration values. This problem can occur when you use a command to change the shared memory value of only one of these two components.</p> <p>To resolve the issue, specify the following shared memory value in the configuration file:</p> <pre>SharedMemoryEnabled=1</pre> <p>Then, restart both components. Do not change the values of the other two shared memory parameters.</p> <p>The <code>SharedMemoryEnabled</code> parameter is stored in the following file:</p> <pre>storage_path\etc\puredisk\agent.cfg</pre>
<p>media manager - system error occurred (174)</p>	<p>If the job details also includes errors similar to the following, it indicates that an image clean-up job failed:</p> <pre>Critical bpdm (pid=610364) sts_delete_image failed: error 2060018 file not found Critical bpdm (pid=610364) image delete failed: error 2060018: file not found</pre> <p>This error occurs if a deduplication backup job fails after the job writes part of the backup to the Media Server Deduplication Pool. NetBackup starts an image cleanup job, but that job fails because the data necessary to complete the image clean-up was not written to the Media Server Deduplication Pool.</p> <p>Deduplication queue processing cleans up the image objects, so you do not need to take corrective action. However, examine the job logs and the deduplication logs to determine why the backup job failed.</p> <p>See “About deduplication queue processing” on page 174.</p>

Client deduplication fails

NetBackup client-side agents (including client deduplication) depend on reverse host name look up of NetBackup server names. Conversely, regular backups depend on forward host name resolution. Therefore, the backup of a client that deduplicates its own data may fail, while a normal backup of the client may succeed.

If a client-side deduplication backup fails, verify that your Domain Name Server includes all permutations of the storage server name.

Also, Symantec recommends that you use fully-qualified domain names for your NetBackup environment.

See “[Use fully qualified domain names](#)” on page 54.

Volume state changes to DOWN when volume is unmounted

If a volume becomes unmounted, NetBackup changes the volume state to DOWN. NetBackup jobs that require that volume fail.

To determine the volume state

- ◆ Invoke the following command on the master server or the media server that functions as the deduplication storage server:

UNIX: `/usr/openv/netbackup/bin/admincmd/nbdevquery -listdv -stype PureDisk -U`

Windows: `install_path\NetBackup\bin\admincmd\nbdevquery -listdv -stype PureDisk -U`

The following example output shows that the `DiskPoolVolume` is UP:

```
Disk Pool Name      : PD_Disk_Pool
Disk Type           : PureDisk
Disk Volume Name    : PureDiskVolume
Disk Media ID       : @aaaab
Total Capacity (GB) : 49.98
Free Space (GB)     : 43.66
Use%                : 12
Status              : UP
Flag                : ReadOnWrite
Flag                : AdminUp
Flag                : InternalUp
Num Read Mounts     : 0
Num Write Mounts    : 1
Cur Read Streams   : 0
Cur Write Streams  : 0
```

To change the volume state to UP

- ◆ Mount the file system

After a brief period of time, the volume state changes to UP. No further action is required.

Errors, delayed response, hangs

Insufficient memory or inadequate host capabilities may cause multiple errors, delayed response, and hangs.

See “[About deduplication server requirements](#)” on page 26.

For virtual machines, Symantec recommends that you do the following:

- Set the memory size of each virtual machine to double the physical memory of the host.
- Set the minimum and the maximum values of each virtual machine to the same value (double the physical memory of the host). These memory settings prevent the virtual memory from becoming fragmented on the disk because it does not grow or shrink.

These recommendations may not be the best configuration for every virtual machine. However, Symantec recommends that you try this solution first when troubleshooting performance issues.

Cannot delete a disk pool

If you cannot delete a disk pool that you believe contains no valid backup images, the following information may help you troubleshoot the problem.

Expired fragments remain on disk

Under some circumstances, the fragments that compose an expired backup image may remain on disk even though the images have expired. For example, if the storage server crashes, normal clean-up processes may not run. In those circumstances, you cannot delete a disk pool because image fragment records still exist. The error message may be similar to the following:

```
DSM has found that one or more volumes in the disk pool diskpoolname has image fragments.
```

To delete the disk pool, you must first delete the image fragments. The `nbdelete` command deletes expired image fragments from disk volumes.

To delete the fragments of expired images

- ◆ Run the following command on the master server:

```
UNIX: /usr/opensv/netbackup/bin/admincmd/nbdelete -allvolumes -force
```

```
Windows: install_path\NetBackup\bin\admincmd\nbdelete -allvolumes -force
```

The `-allvolumes` option deletes expired image fragments from all volumes that contain them.

The `-force` option removes the database entries of the image fragments even if fragment deletion fails.

Incomplete SLP duplication jobs

Incomplete storage lifecycle policy duplication jobs may prevent disk pool deletion. You can determine if incomplete jobs exist and then cancel them.

To cancel storage lifecycle policy duplication jobs

- 1 Determine if incomplete SLP duplication jobs exist by running the following command on the master server:

```
UNIX: install_path\NetBackup\bin\admincmd\nbstlutil stlilist
-incomplete
```

```
Windows: /usr/opensv/netbackup/bin/admincmd/nbstlutil stlilist
-incomplete
```

- 2 Cancel the incomplete jobs by running the following command for each backup ID returned by the previous command (`xxxxx` represents the backup ID):

```
UNIX: install_path\NetBackup\bin\admincmd\nbstlutil cancel
-backupid xxxxx
```

```
Windows: /usr/opensv/netbackup/bin/admincmd/nbstlutil cancel
-backupid xxxxx
```

Media open error (83)

To diagnose and resolve media open errors that may occur during backups, see the following possible causes and corrective actions:

Possible causes	The NetBackup Deduplication Engine (<code>spoold</code>) was too busy to respond to the deduplication process in a timely manner.
	The NetBackup Deduplication Manager (<code>spad</code>) was too busy to respond to the deduplication process in a timely manner.

Diagnosis Examine what caused the core media server deduplication processes (`spad` and `spoold`) to be unresponsive. Were they temporarily busy (such as queue processing in progress)? Do too many jobs run concurrently?

See [“About deduplication performance”](#) on page 52.

Status 83 is a generic error for the duplications. The NetBackup `bpdm` log provides additional information for determining the specific issue.

Media write error (84)

To diagnose and resolve media write errors that may occur during backups, see the following possible causes and corrective actions:

The NetBackup Deduplication Engine (<code>spoold</code>) was too busy to respond.	Examine the Disk Logs report for PureDisk errors. Examine the disk monitoring services log files for details from the deduplication plug-in. See “Viewing disk reports” on page 145.
The NetBackup Deduplication Engine rebuilt the deduplication cache when the backup occurred.	Restarting the NetBackup Deduplication Engine causes the cache to be rebuilt. During the cache rebuild no backups are accepted.
Data removal is running.	Data cannot be backed up at the same time as it is removed. See “About deduplication queue processing” on page 174.
A user tampered with the storage.	Users must not add files to, change files on, or delete files from the storage. If a file was added, remove it.
Storage capacity was increased.	If you grew the storage, you must restart the NetBackup services on the storage server so the new capacity is recognized.
The storage is full.	If possible, increase the storage capacity. See “About adding additional storage” on page 70.
The deduplication pool is down.	Change the state to up. See “Changing the deduplication disk pool state” on page 162.
Firewall ports are not open.	Ensure that ports 10082 and 10102 are open in any firewalls between the deduplication hosts.

Host name resolution problems.

Client-side deduplication can fail if the client cannot resolve the host name of the server. More specifically, the error can occur if the storage server was configured with a short name and the client tries to resolve a fully qualified domain name.

To determine which name the client uses for the storage server, examine the deduplication host configuration file on the client.

See [“About the deduplication host configuration file”](#) on page 132.

To fix this problem, configure your network environment so that all permutations of the storage server name resolve.

Symantec recommends that you use fully qualified domain names.

See [“Use fully qualified domain names”](#) on page 54.

Storage full conditions

Operating system tools such as the UNIX `df` command do not report deduplication disk usage accurately. The operating system commands may report that the storage is full when it is not. NetBackup tools let you monitor storage capacity and usage more accurately.

See [“About deduplication storage capacity and usage reporting”](#) on page 141.

See [“About deduplication container files”](#) on page 143.

See [“Viewing storage usage within deduplication container files”](#) on page 143.

Examining the disk log reports for threshold warnings can give you an idea of when a storage full condition may occur.

How NetBackup performs maintenance can affect when storage is freed up for use.

See [“About deduplication queue processing”](#) on page 174.

See [“Data removal process”](#) on page 224.

Although not advised, you can reclaim free space manually.

See [“Processing the deduplication transaction queue manually”](#) on page 174.

Viewing disk errors and events

You can view disk errors and events in several ways, as follows:

- The Disk Logs report.
See “Viewing disk reports” on page 145.
- The NetBackup `bpererror` command with the `-disk` option reports on disk errors. The command resides in the following directories:
UNIX: `/usr/opensv/netbackup/bin/admincmd`
Windows: `install_path\Veritas\NetBackup\bin\admincmd`

Deduplication event codes and messages

The following table shows the deduplication event codes and their messages. Event codes appear in the `bpererror` command `-disk` output and in the disk reports in the NetBackup Administration Console.

Table 8-4 Deduplication event codes and messages

Event #	Event Severity	NetBackup Severity	Message example
1000	2	Error	Operation configload/reload failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1001	2	Error	Operation configload/reload failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1002	4	Warning	The open file limit exceeded in server PureDisk:server1.symantecs.org on host server1.symantecs.org. Will attempt to continue further.
1003	2	Error	A connection request was denied on the server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1004	1	Critical	Network failure occurred in server PureDisk:server1.symantecs.org on host server1.symantecs.org.

Table 8-4 Deduplication event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1013	1	Critical	Task session start request on server PureDisk:server1.symantecs.org on host server1.symantecs.org got an unexpected error.
1008	2	Error	Task Aborted; An unexpected error occurred during communication with remote system in server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1009	8	Authorization	Authorization request from <IP> for user <USER> denied (<REASON>).
1010	2	Error	Task initialization on server PureDisk:server1.symantecs.org on host server1.symantecs.org got an unexpected error.
1011	2	Error	Task ended on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1012	2	Error	A request for agent task was denied on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1014	1	Critical	Task session start request on server PureDisk:server1.symantecs.org on host server1.symantecs.org got an unexpected error.
1015	1	Critical	Task creation failed, could not initialize task class on server PureDisk:server1.symantecs.org on host server1.symantecs.org.

Table 8-4 Deduplication event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1017	1	Critical	Service Symantec DeduplicationEngine exit on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Please check the server log for the probable cause of this error. The application has terminated.
1018	16	Info	Startup of Symantec Deduplication Engine completed successfully on server1.symantecs.org.
1019	1	Critical	Service Symantec DeduplicationEngine restart on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Please check the server log for the probable cause of this error. The application has restarted.
1020	1	Critical	Service Symantec Deduplication Engine connection manager restart failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Please check the server log for the probable cause of this error.The application has failed to restart.
1028	1	Critical	Service Symantec DeduplicationEngine abort on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Please check the server log for the probable cause of this error.The application has caught an unexpected signal.
1029	1	Critical	Double backend initialization failure; Could not initialize storage backend or cache failure detected on host PureDisk:server1.symantecs.org in server server1.symantecs.org.

Table 8-4 Deduplication event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1030	1	Critical	Operation Storage Database Initialization failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1031	1	Critical	Operation Content router context initialization failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1032	1	Critical	Operation log path creation/print failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1036	4	Warning	Operation a transaction failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1037	4	Warning	Transaction failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Transaction will be retried.
1040	2	Error	Operation Database recovery failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1043	2	Error	Operation Storage recovery failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1044	multiple	multiple	The usage of one or more system resources has exceeded a warning level. Operations will or could be suspended. Please take action immediately to remedy this situation.
1047	2	Error	CRC mismatch detected; possible corruption in server PureDisk:server1.symantecs.org on host server1.symantecs.org.

Table 8-4 Deduplication event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1057			A data corruption has been detected. The data consistency check detected a data loss or data corruption in the Media Server Deduplication Pool (MSDP) and reported the affected backups. Search <code>storaged.log</code> on the server for the affected backups and contact technical support.
1058			A data inconsistency has been detected and corrected automatically. Explanation: The data consistency check detected a potential data loss and fixed it automatically in the Media Server Deduplication Pool (MSDP). Search the <code>storaged.log</code> file on the pertinent media server. Contact support to investigate the root cause if the problem persists.
2000		Error	Low space threshold exceeded on the partition containing the storage database on server PureDisk:server1.symantecs.org on host server1.symantecs.org.

Host replacement, recovery, and uninstallation

This chapter includes the following topics:

- [Replacing the deduplication storage server host computer](#)
- [Recovering from a deduplication storage server disk failure](#)
- [Recovering from a deduplication storage server failure](#)
- [Recovering the storage server after NetBackup catalog recovery](#)
- [About uninstalling media server deduplication](#)
- [Removing media server deduplication](#)

Replacing the deduplication storage server host computer

If you replace the deduplication storage server host computer, use these instructions to install NetBackup and reconfigure the deduplication storage server. The new host cannot host a deduplication storage server already.

Reasons to replace the host include a lease swap or perhaps the current deduplication storage server host does not meet your performance requirements.

When you configure the new host, you can do the following:

- Use the same host name or a different name.
- Use the same network interface (if the original server used a specific network interface) or a different network interface. Alternatively, you do not have to use a specific network interface.

- Use the same storage path or a different storage path. If you use a different storage path, you must move the deduplication storage to that new location.

Warning: The new host must use the same operating system and the same byte order as the old host. If it does not, you cannot access the deduplicated data.

In computing, endianness describes the byte order that represents data: big endian and little endian. For example, SPARC processors and Intel processors use different byte orders. Therefore, you cannot replace an Oracle Solaris SPARC host with an Oracle Solaris host that has an Intel processor.

Table 9-1 Replacing a deduplication storage server host computer

Step	Task	Procedure
Step 1	Expire the backup images	Expire all backup images that reside on the deduplication disk storage. Warning: Do not delete the images. They are imported back into NetBackup later in this process. If you use the <code>bpxpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter. See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .
Step 2	Delete the storage units that use the disk pool	See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .
Step 3	Delete the disk pool	See “Deleting a deduplication disk pool” on page 172.
Step 4	Delete the deduplication storage server	See “Deleting a deduplication storage server” on page 158.
Step 5	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See “Deleting a deduplication host configuration file” on page 132.
Step 6	Delete the credentials on deduplication servers	If you have load balancing servers, delete the NetBackup Deduplication Engine credentials on those media servers. See “Deleting credentials from a load balancing server” on page 161.

Table 9-1 Replacing a deduplication storage server host computer (*continued*)

Step	Task	Procedure
Step 7	Configure the new host so it meets deduplication requirements	<p>When you configure the new host, you can do the following:</p> <ul style="list-style-type: none"> ■ Use the same host name or a different name. ■ Use the same network interface (if the original server used a specific network interface) or a different network interface. Alternatively, you do not have to use a specific network interface. ■ Use the same storage path or a different storage path. If you use a different storage path, you must move the deduplication storage to that new location. <p>See “About NetBackup deduplication servers” on page 25. See “About deduplication server requirements” on page 26.</p>
Step 8	Connect the storage to the host	<p>Use the storage path that you configured for this replacement host. See the computer or the storage vendor's documentation.</p>
Step 9	Install the NetBackup media server software on the new host	<p>See the <i>NetBackup Installation Guide for UNIX and Linux</i>. See the <i>NetBackup Installation Guide for Windows</i>.</p>
Step 10	Reconfigure deduplication	<p>See “Configuring NetBackup media server deduplication” on page 76.</p>
Step 12	Import the backup images	<p>See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i>.</p>

Recovering from a deduplication storage server disk failure

If recovery mechanisms do not protect the disk on which the NetBackup software resides, the deduplication storage server configuration is lost if the disk fails. This topic describes how to recover from a system disk or program disk failure where the disk was not backed up.

Note: This procedure describes recovery of the disk on which the NetBackup media server software resides not the disk on which the deduplicated data resides. The disk may or may not be the system boot disk.

After recovery, your NetBackup deduplication environment should function normally. Any valid backup images on the deduplication storage should be available for restores.

Symantec recommends that you use NetBackup to protect the deduplication storage server system or program disks. You then can use NetBackup to restore that media server if the disk on which NetBackup resides fails and you have to replace it.

Table 9-2 Process to recover from media server disk failure

Step	Task	Procedure
Step 1	Replace the disk.	If the disk is a system boot disk, also install the operating system. See the hardware vendor and operating system documentation.
Step 2	Mount the storage.	Ensure that the storage and database are mounted at the same locations. See the storage vendor's documentation.
Step 3	Install and license the NetBackup media server software.	See the <i>NetBackup Installation Guide for UNIX and Linux</i> . See the <i>NetBackup Installation Guide for Windows</i> . See “About the deduplication license key” on page 72.
Step 4	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See “Deleting a deduplication host configuration file” on page 132.
Step 5	Delete the credentials on deduplication servers	If you have load balancing servers, delete the NetBackup Deduplication Engine credentials on those media servers. See “Deleting credentials from a load balancing server” on page 161.
Step 6	Add the credentials to the storage server	Add the NetBackup Deduplication Engine credentials to the storage server. See “Adding NetBackup Deduplication Engine credentials” on page 160.
Step 7	Get a configuration file template	If you did not save a storage server configuration file before the disk failure, get a template configuration file. See “Saving the deduplication storage server configuration” on page 129.
Step 8	Edit the configuration file	See “Editing a deduplication storage server configuration file” on page 129.
Step 9	Configure the storage server	Configure the storage server by uploading the configuration from the file you edited. See “Setting the deduplication storage server configuration” on page 131.
Step 10	Add load balancing servers	If you use load balancing servers in your environment, add them to your configuration. See “Adding a deduplication load balancing server” on page 117.

Recovering from a deduplication storage server failure

To recover from a permanent failure of the storage server host computer, use the process that is described in this topic.

When you configure the new host, you can do the following:

- Use the same host name or a different name.
- Use the same network interface (if the original server used a specific network interface) or a different network interface. Alternatively, you do not have to use a specific network interface.
- Use the same storage path or a different storage path. If you use a different storage path, you must move the deduplication storage to that new location.

Warning: The new host must use the same operating system and the same byte order as the old host. If it does not, you cannot access the deduplicated data.

In computing, endianness describes the byte order that represents data: big endian and little endian. For example, SPARC processors and Intel processors use different byte orders. Therefore, you cannot replace an Oracle Solaris SPARC host with an Oracle Solaris host that has an Intel processor.

Table 9-3 Recover from a deduplication storage server failure

Step	Task	Procedure
Step 1	Expire the backup images	<p>Expire all backup images that reside on the deduplication disk storage.</p> <p>Warning: Do not delete the images. They are imported back into NetBackup later in this process.</p> <p>If you use the <code>bpxpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter.</p> <p>See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i></p> <p>See the <i>NetBackup Administrator's Guide for Windows, Volume I</i>.</p>
Step 2	Delete the storage units that use the disk pool	<p>See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i></p> <p>See the <i>NetBackup Administrator's Guide for Windows, Volume I</i>.</p>
Step 3	Delete the disk pool	See “Deleting a deduplication disk pool” on page 172.
Step 4	Delete the deduplication storage server	See “Deleting a deduplication storage server” on page 158.

Table 9-3 Recover from a deduplication storage server failure (*continued*)

Step	Task	Procedure
Step 5	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See “Deleting a deduplication host configuration file” on page 132.
Step 6	Delete the credentials on deduplication servers	If you have load balancing servers, delete the NetBackup Deduplication Engine credentials on those media servers. See “Deleting credentials from a load balancing server” on page 161.
Step 7	Configure the new host so it meets deduplication requirements	When you configure the new host, you can do the following: <ul style="list-style-type: none"> ■ Use the same host name or a different name. ■ Use the same network interface (if the original server used a specific network interface) or a different network interface. Alternatively, you do not have to use a specific network interface. ■ Use the same storage path or a different storage path. If you use a different storage path, you must make the storage available over that storage path. See “About NetBackup deduplication servers” on page 25. See “About deduplication server requirements” on page 26.
Step 8	Connect the storage to the host	Use the storage path that you configured for this replacement host. See the computer or the storage vendor's documentation.
Step 9	Install the NetBackup media server software on the new host	See the <i>NetBackup Installation Guide for UNIX and Linux</i> . See the <i>NetBackup Installation Guide for Windows</i> .
Step 10	Reconfigure deduplication	You must use the same credentials for the NetBackup Deduplication Engine. See “Configuring NetBackup media server deduplication” on page 76.
Step 11	Import the backup images	See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .

Recovering the storage server after NetBackup catalog recovery

If a disaster requires a recovery of the NetBackup catalog, you must set the storage server configuration after the NetBackup catalog is recovered.

See [“Setting the deduplication storage server configuration”](#) on page 131.

Symantec recommends that you save your storage server configuration.

See [“Save the deduplication storage server configuration”](#) on page 59.

Information about recovering the master server is available.

See the *NetBackup Troubleshooting Guide*.

About uninstalling media server deduplication

You cannot uninstall media server deduplication components separately from NetBackup. The deduplication components are installed when you install NetBackup software, and they are uninstalled when you uninstall NetBackup software.

Other topics describe related procedures, as follow:

- Reconfigure an existing deduplication environment.
 See [“Changing the deduplication storage server name or storage path”](#) on page 155.
- Deactivate deduplication and remove the configuration files and the storage files.
 See [“Removing media server deduplication”](#) on page 213.

Removing media server deduplication

You cannot remove the deduplication components from a NetBackup media server. You can disable the components and remove the deduplication storage files and the catalog files. The host remains a NetBackup media server.

This process assumes that all backup images that reside on the deduplication disk storage have expired.

Warning: If you remove deduplication and valid NetBackup images reside on the deduplication storage, data loss may occur.

Table 9-4 Remove media server deduplication

Step	Task	Procedure
Step 1	Remove client deduplication	Remove the clients that deduplicate their own data from the client deduplication list. See “Disabling client-side deduplication for a client” on page 173.

Table 9-4 Remove media server deduplication (*continued*)

Step	Task	Procedure
Step 2	Delete the storage units that use the disk pool	See the <i>NetBackup Administrator's Guide for UNIX and Linux, Volume I</i> See the <i>NetBackup Administrator's Guide for Windows, Volume I</i> .
Step 3	Delete the disk pool	See “Deleting a deduplication disk pool” on page 172.
Step 4	Delete the deduplication storage server	See “Deleting a deduplication storage server” on page 158. Deleting the deduplication storage server does not alter the contents of the storage on physical disk. To protect against inadvertent data loss, NetBackup does not automatically delete the storage when you delete the storage server.
Step 5	Delete the configuration	Delete the deduplication configuration. See “Deleting the deduplication storage server configuration” on page 159.
Step 6	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See “Deleting a deduplication host configuration file” on page 132.
Step 7	Delete the storage directory and the database directory	Delete the storage directory and database directory. (Using a separate database directory was an option when you configured deduplication.) Warning: If you delete the storage directory and valid NetBackup images reside on the deduplication storage, data loss may occur. See the operating system documentation.

Deduplication architecture

This chapter includes the following topics:

- [Deduplication storage server components](#)
- [Media server deduplication process](#)
- [Deduplication client components](#)
- [Client-side deduplication backup process](#)
- [About deduplication fingerprinting](#)
- [Data removal process](#)

Deduplication storage server components

[Figure 10-1](#) is a diagram of the storage server components.

Figure 10-1 Storage server deduplication components

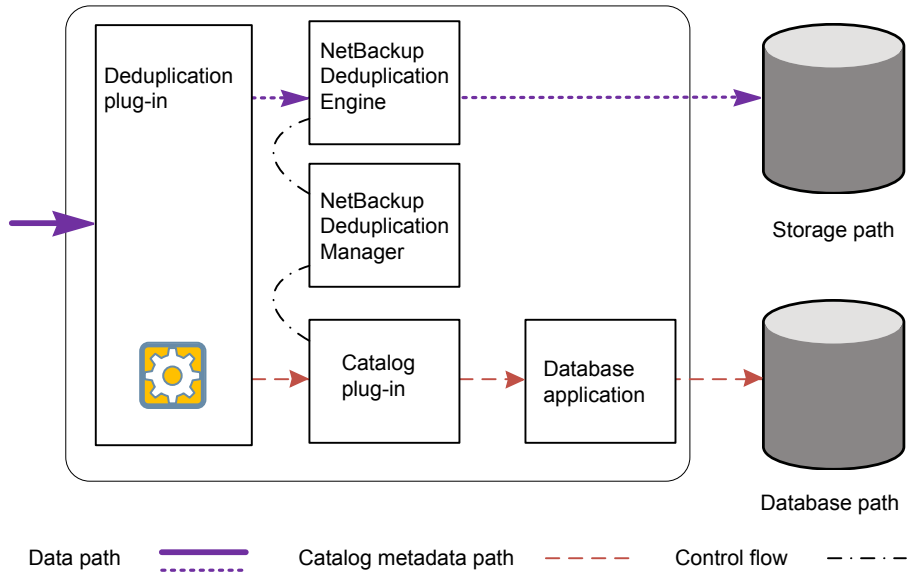


Table 10-1 describes the components.

Table 10-1 NetBackup deduplication components

Component	Description
Deduplication plug-in	<p>The deduplication plug-in is the data interface to the NetBackup Deduplication Engine on the storage server. The deduplication plug-in does the following:</p> <ul style="list-style-type: none"> ■ Separates the file's metadata from the file's content. ■ Deduplicates the content (separates files into segments). ■ Controls the data stream from NetBackup to the NetBackup Deduplication Engine and vice versa. <p>The plug-in runs on the deduplication storage server. The plug-in also runs on load balancing servers and on the clients that deduplicate their own data.</p>
NetBackup Deduplication Engine	<p>The NetBackup Deduplication Engine is one of the storage server core components. It stores and manages deduplicated file data.</p> <p>The binary file name is <code>spoold</code>, which is short for storage pool daemon; do not confuse it with a print spooler daemon. The <code>spoold</code> process appears as the NetBackup Deduplication Engine in the NetBackup Administration Console.</p>

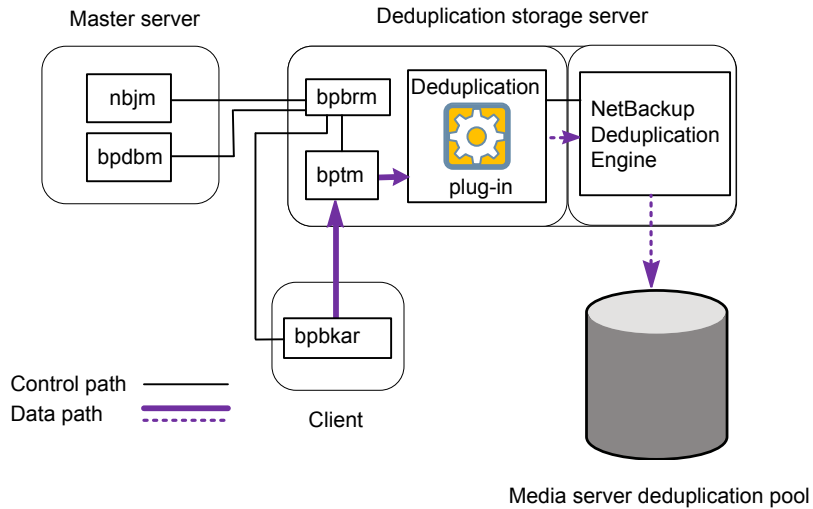
Table 10-1 NetBackup deduplication components (*continued*)

Component	Description
NetBackup Deduplication Manager	<p>The deduplication manager is one of the storage server core components. The deduplication manager maintains the configuration and controls internal processes, optimized duplication, security, and event escalation.</p> <p>The deduplication manager binary file name is <code>spad</code>. The <code>spad</code> process appears as the NetBackup Deduplication Manager in the NetBackup Administration Console.</p>
Catalog plug-in	The catalog plug-in implements a standardized catalog API, which lets the NetBackup Deduplication Engine communicate with the back-end database process. The catalog plug-in translates deduplication engine catalog calls into the calls that are native to the back-end database application.
Database application	The database application communicates with the catalog plug-in. The database application writes data to and reads data from the database.
Deduplication database	<p>The deduplication database stores and manages the metadata of deduplicated files. The metadata includes a unique fingerprint that identifies the file's content. The metadata also includes information about the file such as its owner, where it resides on a client, when it was created, and other information.</p> <p>NetBackup uses the PostgreSQL database for the deduplication database.</p> <p>You can use the NetBackup <code>bpps</code> command to view the database process (<code>postgres</code>).</p> <p>The deduplication database is separate from the NetBackup catalog. The NetBackup catalog maintains the usual NetBackup backup image information.</p>

Media server deduplication process

[Figure 10-2](#) shows the backup process when a media server deduplicates the backups. The destination is a **Media Server Deduplication Pool**. A description follows.

Figure 10-2 Media server deduplication process



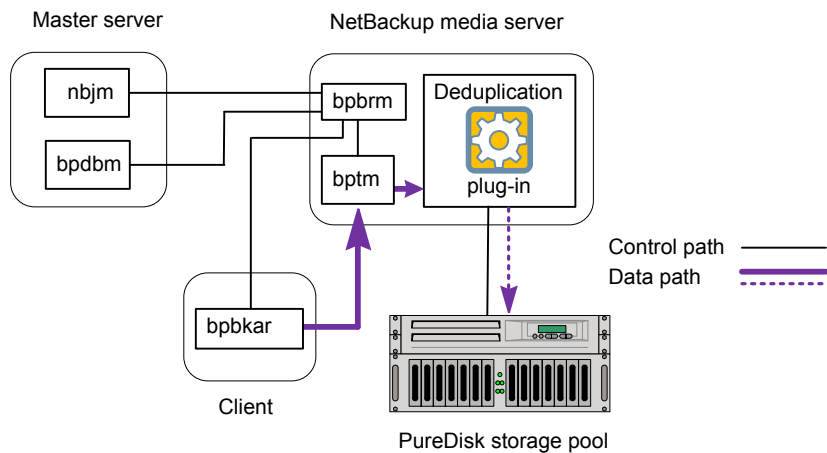
The following list describes the backup process when a media server deduplicates the backups and the destination is a **Media Server Deduplication Pool**:

- The NetBackup Job Manager (nbjm) starts the Backup/Restore Manager (bpbm) on a media server.
- The Backup/Restore Manager starts the bpbm process on the media server and the bpbkar process on the client.
- The Backup/Archive Manager (bpbkar) on the client generates the backup images and moves them to the media server bpbm process. The Backup/Archive Manager also sends the information about files within the image to the Backup/Restore Manager (bpbm). The Backup/Restore Manager sends the file information to the bpbm process on the master server for the NetBackup database.
- The bpbm process moves the data to the deduplication plug-in.
- The deduplication plug-in retrieves a list of fingerprints from the last full backup for the client from the NetBackup Deduplication Engine. The list is used as a cache so the plug-in does not have to request each fingerprint from the engine.
- The deduplication plug-in performs file fingerprinting calculations.
- The deduplication plug-in compares the file fingerprints and the segment fingerprints against the fingerprint list in its cache.

- The deduplication plug-in sends only unique data segments to the NetBackup Deduplication Engine on the storage server. The NetBackup Deduplication Engine writes the data to the **Media Server Deduplication Pool**.

Figure 10-3 shows the backup process when a media server deduplicates the backups. The destination is a **PureDisk Deduplication Pool**. A description follows.

Figure 10-3 Media server deduplication process to a PureDisk storage pool



The following list describes the backup process when a media server deduplicates the backups and the destination is a **PureDisk Deduplication Pool**:

- The NetBackup Job Manager (`nbjm`) starts the Backup/Restore Manager (`bpbrm`) on a media server.
- The Backup/Restore Manager starts the `bptm` process on the media server and the `bpbkar` process on the client).
- The Backup/Archive Manager (`bpbkar`) generates the backup images and moves them to the media server `bptm` process. The Backup/Archive Manager also sends the information about files within the image to the Backup/Restore Manager (`bpbrm`). The Backup/Restore Manager sends the file information to the `bpdbm` process on the master server for the NetBackup database.
- The `bptm` process moves the data to the deduplication plug-in.
- The deduplication plug-in retrieves a list of fingerprints from the last full backup for the client from the PureDisk storage pool authority. The list is used as a cache so the plug-in does not have to request each fingerprint individually.

- The deduplication plug-in compares the file fingerprints and the segment fingerprints against the fingerprint list in its cache.
- The deduplication plug-in performs file fingerprinting calculations.
- The deduplication plug-in sends only unique data segments to the **PureDisk Deduplication Pool**.

Deduplication client components

Table 10-2 describes the client deduplication components.

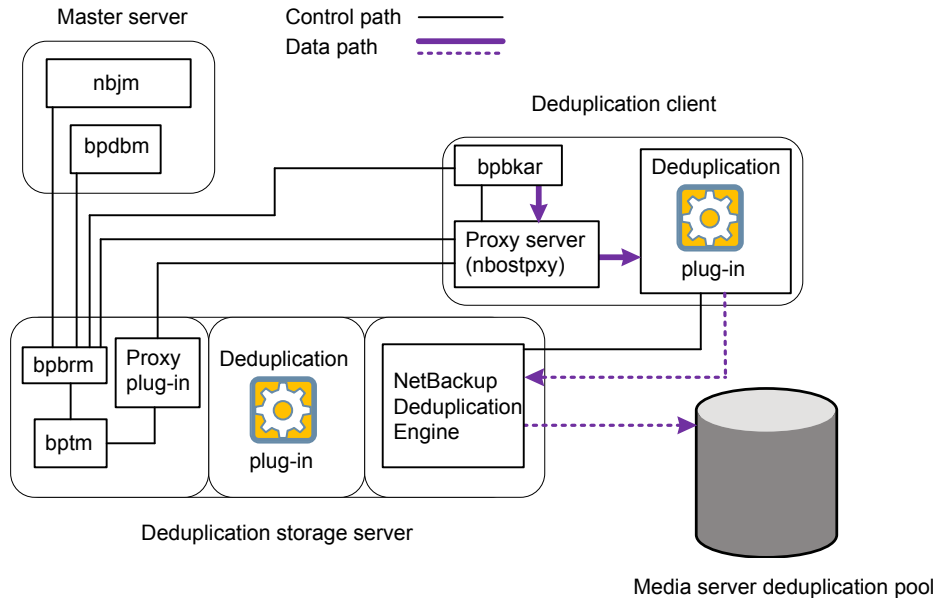
Table 10-2 Client deduplication components

Component	Host	Description
Deduplication plug-in	Client	<p>The deduplication plug-in is the data interface to the NetBackup Deduplication Engine on the deduplication storage server. The deduplication plug-in does the following:</p> <ul style="list-style-type: none"> ■ Separates the file’s metadata from the file’s content. ■ Deduplicates the content (separates files into segments). ■ Controls the data stream from NetBackup to the NetBackup Deduplication Engine and vice versa.
Proxy server	Client	The OpenStorage proxy server (<code>nbostrpxy</code>) manages control communication with the media server.
Proxy plugin	Media server	The proxy plug-in manages control communication with the client.

Client–side deduplication backup process

Figure 10-4 shows the backup process of a client that deduplicates its own data. The destination is a media server deduplication pool. A description follows.

Figure 10-4 Deduplication client backup to a media server deduplication pool



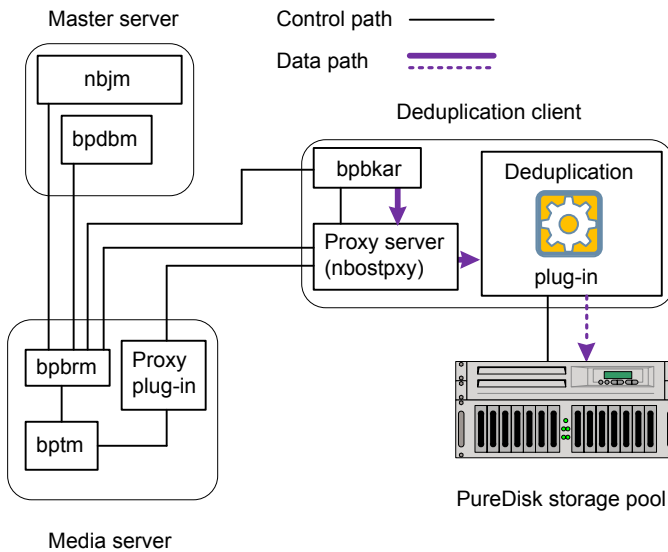
The following list describes the backup process for a deduplication client to a media server deduplication pool:

- The NetBackup Job Manager (`nbjm`) starts the Backup/Restore Manager (`bpbrm`) on a media server.
- The Backup/Restore Manager probes the client to determine if it is configured and ready for deduplication.
- If the client is ready, the Backup/Restore Manager starts the following processes: The OpenStorage proxy server (`nbostpxy`) on the client and the data moving processes (`bpbkar`) on the client and `bptm` on the media server). NetBackup uses the proxy plug-in on the media server to route control information from `bptm` to `nbostpxy`.
- The Backup/Archive Manager (`bpbkar`) generates the backup images and moves them to the client `nbostpxy` process by shared memory. The Backup/Archive Manager also sends the information about files within the image to the Backup/Restore Manager (`bpbrm`). The Backup/Restore Manager sends the file information to the `bpdbm` process on the master server for the NetBackup database.
- The client `nbostpxy` process moves the data to the deduplication plug-in.

- The deduplication plug-in retrieves a list of fingerprints from the last full backup for the client from the NetBackup Deduplication Engine. The list is used as a cache so the plug-in does not have to request each fingerprint from the engine.
- The deduplication plug-in performs file fingerprinting calculations.
- The deduplication plug-in sends only unique data segments to the storage server, which writes the data to the media server deduplication pool.

Figure 10-5 shows the backup process of a client that deduplicates its own data. The destination is a PureDisk storage pool. A description follows.

Figure 10-5 Deduplication client backup to a PureDisk storage pool



The following list describes the backup process for a deduplication client to a media server deduplication pool:

- The NetBackup Job Manager (nbjm) starts the Backup/Restore Manager (bpbm) on a media server.
- The Backup / Restore Manager probes the client to determine if it is configured and ready for deduplication.
- If the client is ready, the Backup/Restore Manager starts the following processes: The OpenStorage proxy server (nbostpxy) on the client and the data moving processes (bpbkar on the client and bptm on the media server).

NetBackup uses the proxy plug-in on the media server to route control information from `bptm` to `nboostpxy`.

- The Backup/Archive Manager (`bpbkar`) generates the backup images and moves them to the client `nboostpxy` process by shared memory. The Backup/Archive Manager also sends the information about files within the image to the Backup/Restore Manager (`bpbrm`). The Backup/Restore Manager sends the file information to the `bpdbm` process on the master server for the NetBackup database.
- The client `nboostpxy` process moves the data to the deduplication plug-in.
- The deduplication plug-in retrieves a list of fingerprints from the last full backup for the client from the NetBackup Deduplication Engine. The list is used as a cache so the plug-in does not have to request each fingerprint from the engine.
- The deduplication plug-in performs file fingerprinting calculations.
- The deduplication plug-in sends only unique data segments to the PureDisk storage pool.

About deduplication fingerprinting

The NetBackup Deduplication Engine uses a unique identifier to identify each file and each file segment that is backed up. The engine identifies files inside the backup images and then processes the files.

The process is known as fingerprinting.

For the first deduplicated backup, the following is the process:

- The deduplication plug-in reads the backup image and separates the image into files.
- The plug-in separates files into segments.
- For each segment, the plug-in calculates the hash key (or fingerprint) that identifies each data segment. To create a hash, every byte of data in the segment is read and added to the hash.
- The plug-in compares its calculated fingerprints to the fingerprints that the NetBackup Deduplication Engine stores on the media server. Two segments that have the same fingerprint are duplicates of each other.
- The plug-in sends unique segments to the deduplication engine to be stored. A unique segment is one for which a matching fingerprint does not exist in the engine already.

The first backup may have a 0% deduplication rate; however, a 0% deduplication rate is unlikely. Zero percent means that all file segments in the backup data are unique.

- The NetBackup Deduplication Engine saves the fingerprint information for that backup.

For subsequent backups, the following is the process:

- The deduplication plug-in retrieves a list of fingerprints from the last full backup for the client from the NetBackup Deduplication Engine. The list is used as a cache so the plug-in does not have to request each fingerprint from the engine.
- The deduplication plug-in reads the backup image and separates the image into files.
- The deduplication plug-in separates files into segments and calculates the fingerprint for each file and segment.
- The plug-in compares each fingerprint against the local fingerprint cache.
- If the fingerprint is not known in the cache, the plug-in requests that the engine verify if the fingerprint already exists. If the fingerprint does not exist, the segment is sent to the engine. If the fingerprint exists, the segment is not sent.

The fingerprint calculations are based on the MD5 algorithm. However, any segments that have different content but the same MD5 hash key get different fingerprints. So NetBackup prevents MD5 collisions.

Data removal process

The following list describes the data removal process for expired backup images:

- NetBackup removes the image record from the NetBackup catalog. NetBackup directs the NetBackup Deduplication Manager to remove the image.
- The deduplication manager immediately removes the image entry and adds a removal request for the image to the database transaction queue. From this point on, the image is no longer accessible.
- When the queue is next processed, the NetBackup Deduplication Engine executes the removal request. The engine also generates removal requests for underlying data segments
- At the successive queue processing, the NetBackup Deduplication Engine executes the removal requests for the segments.

Storage is reclaimed after two queue processing runs; that is, in one day. However, data segments of the removed image may still be in use by other images.

If you manually delete an image that has expired within the previous 24 hours, the data becomes garbage. It remains on disk until removed by the next garbage collection process.

See [“About deduplication queue processing”](#) on page 174.

See [“Deleting backup images”](#) on page 173.

NetBackup appliance deduplication

This appendix includes the following topics:

- [About NetBackup appliance deduplication](#)
- [About Fibre Channel to a NetBackup 5020 appliance](#)
- [Enabling Fibre Channel to a NetBackup 5020 appliance](#)
- [Disabling Fibre Channel to a NetBackup 5020 appliance](#)
- [Displaying NetBackup 5020 appliance Fibre Channel port information](#)

About NetBackup appliance deduplication

NetBackup appliances are hardware and software solutions from Symantec that combine a host and storage with the Symantec backup software. The appliances offer customers easy and convenient deployment options for Symantec's industry-leading backup and deduplication technologies. The appliances enable efficient, storage-optimized data protection for the datacenter, remote office, and virtual environments.

Symantec's NetBackup appliance family consists of the following two series:

- NetBackup 5200 series of enterprise backup appliances that is based on the NetBackup backup platform. The 5200 series provides between 4 TB to 32 TB of deduplication storage.

A NetBackup 5200 series appliance can be a destination for optimized duplication from a NetBackup **Media Server Deduplication Pool**.

- NetBackup 5000 series of scalable deduplication appliances that is based on the PureDisk backup platform. The 5000 series is scalable from 16 TB to 192 TB of storage.
A NetBackup 5000 series can be a storage destination for both the NetBackup Client Deduplication Option and the NetBackup Media Server Deduplication Option.

The NetBackup appliances share many common features, as follows:

- Easy to install, configure, and use.
- Modular capacity to fulfill your storage needs.
- A solution for the datacenter, remote office and branch office, and virtual machine backups. Source or target deduplication. Optimized synthetic backup to minimize data movement. Tape support for long-term data retention.
- Built in disk to disk replication for disaster recovery and an alternative solution to tape based vaulting
- Enterprise-class hardware and software. Hardware monitoring with a call home feature.

About Fibre Channel to a NetBackup 5020 appliance

With this release of NetBackup, you can use Fibre Channel for data traffic to a Symantec NetBackup 5020 appliance. The appliance must be at software release 1.3 or later.

The 5020 appliance provides the same functionality as a traditional PureDisk environment, which is a software only environment that is installed on your hardware. This appliance solution can function as the PureDisk storage destination for the following operations:

- A target for NetBackup client backups.
- A target for optimized duplication from a NetBackup **Media Server Deduplication Pool**.
- A source for optimized duplication to another deduplication destination.

For any operation that involves NetBackup PureDisk in this guide, PureDisk means both a traditional PureDisk environment and a NetBackup 5020 appliance.

Requirements for the NetBackup media server:

- NBU 7.5.
- A supported operating system.

See the NetBackup operating system compatibility list on the [NetBackup landing page](#) on the Symantec support Web site.

- One Qlogic 2562 (ISP 2532) HBA with two ports.

Limitations for this solution:

- The appliance supports a maximum of 200 concurrent backup jobs.

Command traffic travels over the IP network. If no Fibre Channel connection is available, backup data travels over the IP network.

For each job, the job details show the amount of data that is transferred over Fibre Channel.

See “[Viewing deduplication job details](#)” on page 138.

For information about configuring the appliance and zoning the Fibre Channel SAN, see the appliance documentation.

Enabling Fibre Channel to a NetBackup 5020 appliance

You can enable Fibre Channel communication for data to a NetBackup 5020 appliance.

Before you enable services, ensure that the target ports of the desired NetBackup 5020 appliance are the only target ports in the Fibre Channel zone.

For information about configuring the appliance and zoning the Fibre Channel SAN, see the appliance documentation.

See “[About Fibre Channel to a NetBackup 5020 appliance](#)” on page 228.

To enable Fibre Channel to a NetBackup 5020 appliance

- ◆ As the root user, run the `dedup_fcmanager.sh` script with the `-e` option as in the following example:

```
/usr/opensv/pdde/pdconfigure/scripts/support/dedup_fcmanager.sh -e
```

```
WARNING: Enabling/disabling Fibre Channel transport may require  
spad to be restarted.
```

```
Do you want to continue? [y/n] y
```

```
FC transport enabled
```

Disabling Fibre Channel to a NetBackup 5020 appliance

You can disable Fibre Channel communication for data to a NetBackup 5020 appliance.

See [“About Fibre Channel to a NetBackup 5020 appliance”](#) on page 228.

To disable Fibre Channel to a NetBackup 5020 appliance

- ◆ As the root user, run the `dedup_fcmanager.sh` script with the `-d` option as in the following example:

```
/usr/opensv/pdde/pdconfigure/scripts/support/dedup_fcmanager.sh -d
```

```
WARNING: Enabling/disabling Fibre Channel transport may require  
spad to be restarted.
```

```
Do you want to continue? [y/n] y
```

```
Restarting services
```

```
FC transport disabled
```

Displaying NetBackup 5020 appliance Fibre Channel port information

From a supported NetBackup media server, you can display the information about the target mode ports on a NetBackup 5020 appliance:

- Port information.

See [“To display Fibre Channel port information on a NetBackup 5020 appliance”](#) on page 231.

- Statistics.

See [“To display Fibre Channel statistics on a NetBackup 5000 series appliance”](#) on page 232.

By default, the top port (port number 1) of the FC HBA in the appliance is configured in the target mode.

Before you display the port information, ensure that the target ports of the desired NetBackup 5000 series appliance are the only target ports in the Fibre Channel zone.

See [“About Fibre Channel to a NetBackup 5020 appliance”](#) on page 228.

To display Fibre Channel port information on a NetBackup 5020 appliance

- ◆ As the root user, run the `dedup_fcmanager.sh` script with the `-r` option as in the following example:

```
/usr/opensv/pdde/pdconfigure/scripts/support/dedup_fcmanager.sh -r
```

```
**** Ports ****
Bus ID   Port WWN                Dev Num  Status  Mode           Speed   Remote Ports

06:00.0  21:00:00:24:FF:xx:xx:xx  3        Online  Target (NBU)   8 gbit/s
06:00.1  21:00:00:24:FF:xx:xx:xx  4        Online  Initiator      8 gbit/s
06:00.0  21:00:00:24:FF:xx:xx:xx  5        Online  Target (NBU)   8 gbit/s
06:00.1  21:00:00:24:FF:xx:xx:xx  6        Online  Initiator      8 gbit/s

**** FC Paths ****
Device      Vendor      Host
/dev/sg3    SYMANTEC    192.168.0.2(5020-Gold.symantecs.org)
/dev/sg5    SYMANTEC    192.168.0.3(5020-Silver.symantecs.org)

**** VLAN ****
The result is based on the scan at Sun, Jan 1 00:00:01 CST 2012
/dev/sg3    192.168.0.2
/dev/sg8    192.168.1.2
/dev/sg5    192.168.0.3

**** Fibre Channel Transport ****
Replication over Fibre Channel is disabled
Backup/Restore over Fibre Channel is disabled
```

This output shows the target mode Fibre Channel ports and the hosts to which Fibre Channel traffic can travel.

To display Fibre Channel statistics on a NetBackup 5000 series appliance

- ◆ As the root user, run the `dedup_fcmanager.sh` script with the `-t` option and interval and repeat arguments. The following command example lists the statistics five times with a one second interval between them:

```
usr/openv/pdde/pdconfigure/scripts/support/dedup_fcmanager.sh -t 1 5
```

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	0	0	0	0

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	2823	12702	0	17144

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	2105	9557	0	13070

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	2130	9597	0	13161

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	2108	9632	0	13136

Some versions of some `qla2xxxx` drivers do not provide KB/s output.

Index

A

- about NetBackup deduplication 13
- about NetBackup deduplication options 13
- about the deduplication host configuration file 132
- appliance deduplication 14–15
- attributes
 - clearing deduplication pool 170
 - clearing deduplication storage server 154
 - OptimizedImage 35
 - setting deduplication pool 164
 - setting deduplication storage server 152
 - viewing deduplication pool 163
 - viewing deduplication storage server 151
- Auto Image Replication
 - nbstserv 105
 - overview 47
 - using MSDP 98

B

- backup
 - client deduplication process 220
- big endian 208, 211
- bpstsinfo command 100
- byte order 208, 211

C

- cache hits field of the job details 140
- capacity
 - adding storage 70
- capacity and usage reporting for deduplication 141
- Capacity managed retention type 110
- changing deduplication server hostname 155
- changing the deduplication storage server name and path 155
- CIFS 67
- clearing deduplication pool attributes 170
- client deduplication
 - about 28
 - components 220
 - disabling for a specific client 173

- client deduplication (*continued*)
 - host requirements 30
 - limitations 30
 - sizing the systems 21
- Common Internet File System 67
- compacting container files 145
- compression
 - and deduplication 34
 - pd.conf file setting 121
- configuring a deduplication pool 81
- configuring a deduplication storage server 79
- configuring a deduplication storage unit 83
- configuring deduplication 76, 78
- container files
 - about 143
 - compaction 145
 - viewing capacity within 143
- contentrouter.cfg file
 - about 127
 - parameters for data integrity checking 179
 - RebaseQuota parameter 181
 - RebaseScatterThreshold parameter 181
 - ServerOptions for encryption 88
- CR sent field of the job details 141
- credentials 32
 - adding NetBackup Deduplication Engine 160
 - changing NetBackup Deduplication Engine 161

D

- data classifications
 - in storage lifecycle policy 107
- data integrity checking
 - about deduplication 175
 - configuring behavior for deduplication 176
 - configuring settings for deduplication 178
- data removal process
 - for deduplication 224
- database system error 192
- deactivating media server deduplication 213
- dedup field of the job details 141

deduplication

- about credentials 32
 - about fingerprinting 223
 - about the license key 72
 - adding credentials 160
 - cache hits field of the job details 140
 - capacity and usage reporting 141
 - changing credentials 161
 - client backup process 220
 - compression 34
 - configuration file 119
 - configuring 76, 78
 - configuring optimized synthetic backups 89
 - container files 143
 - CR sent field of the job details 141
 - data removal process 224
 - dedup field of the job details 141
 - encryption 34
 - event codes 202
 - how it works 15
 - license key for 72
 - licensing 72
 - limitations 28
 - media server process 217
 - network interface 33
 - node 26
 - performance 52
 - planning deployment 20
 - requirements for optimized within the same domain 37
 - scaling 55
 - scanned field of the job details 141
 - storage capacity 68
 - storage destination 22
 - storage management 70
 - storage paths 69
 - storage requirements 66
 - storage unit properties 85
 - stream rate field of the job details 141
 - supported systems 21
- deduplication configuration file
- editing 126
 - settings 119
- deduplication data integrity checking
- about 175
 - configuring behavior for 176
 - configuring settings 178
- deduplication database
- about 217

deduplication database *(continued)*

- log file 188
- deduplication deduplication pool. *See* deduplication pool
- deduplication disk volume
- changing the state 171
 - determining the state 171
- deduplication encryption
- enabling 87
- deduplication host configuration file
- about 132
 - deleting 132
- deduplication hosts
- and firewalls 33
 - client requirements 30
 - load balancing server 26
 - server requirements 26
 - storage server 25
- deduplication logs
- about 187
 - client deduplication proxy plug-in log 188
 - client deduplication proxy server log 188
 - configuration script 188
 - deduplication database 188
 - deduplication plug-in log 189
 - NetBackup Deduplication Engine 189
 - NetBackup Deduplication Manager 189
 - VxUL deduplication logs 190
- deduplication node
- about 26
 - adding a load balancing server 117
 - removing a load balancing server 157
- deduplication optimized synthetic backups
- about 35
- deduplication plug-in
- about 216
 - log file 189
- deduplication pool. *See* deduplication pool
- about 79
 - changing properties 165
 - changing the state 162
 - clearing attributes 170
 - configuring 81
 - deleting 172
 - determining the state 162
 - properties 81
 - setting attributes 164
 - viewing 162
 - viewing attributes 163

- deduplication port usage
 - about 33
 - troubleshooting 200
 - deduplication processes do not start 195
 - deduplication rate
 - how file size affects 54
 - monitoring 137
 - deduplication registry
 - resetting 132
 - deduplication servers
 - about 25
 - components 215
 - host requirements 26
 - deduplication storage capacity
 - about 68
 - viewing capacity in container files 143
 - deduplication storage destination 22
 - deduplication storage paths 69
 - deduplication storage requirements 66
 - deduplication storage server
 - about 25
 - change the name 155
 - changing properties 153
 - clearing attributes 154
 - components 215
 - configuration failure 193
 - configuring 79
 - defining target for Auto Image Replication 49
 - deleting 158
 - deleting the configuration 159
 - determining the state 150
 - editing configuration file 129
 - getting the configuration 129
 - recovery 211
 - replacing the host 207
 - setting attributes 152
 - setting the configuration 131
 - viewing 150
 - viewing attributes 151
 - deduplication storage server configuration file
 - about 128
 - deduplication storage server name
 - changing 155
 - deduplication storage type 22
 - Deduplication storage unit
 - Only use the following media servers 85
 - Use any available media server 85
 - deleting backup images 173
 - deleting deduplication host configuration file 132
 - disaster recovery
 - protecting the data 58
 - recovering the storage server after catalog recovery 212
 - disk failure
 - deduplication storage server 209
 - disk logs 146
 - disk logs report 143
 - disk pool
 - cannot delete 198
 - disk pool status report 143, 146
 - disk storage unit report 146
 - Disk type 85
 - disk volume
 - changing the state 171
 - determining the state of a deduplication 171
 - volume state changes to down 197
 - domains
 - replicating backups to another. *See* Auto Image Replication
- ## E
- Enable file recovery from VM backup 111
 - encryption
 - and deduplication 34
 - enabling for deduplication 87
 - pd.conf file setting 122
 - endian
 - big 208, 211
 - little 208, 211
 - event codes
 - deduplication 202
 - Expire after copy retention type 110
- ## F
- file system
 - CIFS 67
 - NFS 67
 - Veritas File System for deduplication storage 70
 - fingerprinting
 - about deduplication 223
 - firewalls and deduplication hosts 33
 - Fixed retention type 110
 - FlashBackup policy
 - Maximum fragment size (storage unit setting) 85
 - FQDN or IP Address property in Resilient Network
 - host properties 114

G

garbage collection. *See* queue processing

H

host requirements 26
how deduplication works 15

I

images on disk report 145
Import operation 104
initial seeding 53
iSCSI 67

L

license information failure
 for deduplication 193
license key
 for deduplication 72
licensing deduplication 72
limitations
 media server deduplication 28
little endian 208, 211
load balancing server
 about 26
 adding to a deduplication node 117
 for deduplication 26
 removing from deduplication node 157
logs
 about deduplication 187
 client deduplication proxy plug-in log 188
 client deduplication proxy server log 188
 deduplication configuration script log 188
 deduplication database log 188
 deduplication plug-in log 189
 disk 146
 NetBackup Deduplication Engine log 189
 NetBackup Deduplication Manager log 189
 VxUL deduplication logs 190

M

maintenance processing. *See* queue processing
Maximum concurrent jobs 86
Maximum fragment size 85
Maximum snapshot limit retention type 110
media server deduplication
 process 217
 sizing the systems 21

Media Server Deduplication Option
 about 23
Media Server Deduplication Pool 98
media server deduplication pool. *See* deduplication pool
migrating from PureDisk to NetBackup deduplication 61
migrating to NetBackup deduplication 62
Mirror retention type 110
MSDP replication
 about 36

N

nbstserv process 105
NetBackup
 naming conventions 22
NetBackup 5000 series appliance
 as a storage destination 23
NetBackup appliance deduplication 14
NetBackup Client Deduplication Option 14
NetBackup deduplication
 about 13
 license key for 72
NetBackup Deduplication Engine
 about 216
 about credentials 32
 adding credentials 160
 changing credentials 161
 logs 189
NetBackup Deduplication Manager
 about 217
 logs 189
NetBackup deduplication options 13
NetBackup Media Server Deduplication Option 14
network interface
 for deduplication 33
NFS 67
node
 deduplication 26

O

OpenStorage appliance deduplication 15
optimized deduplication
 configuring bandwidth 89
optimized deduplication copy
 configuring 94
 guidance for 39
 limitations 38

optimized deduplication copy *(continued)*

- push configuration 41
- separate network for 92

optimized MSDP deduplication

- about the media server in common within the same domain 39
- pull configuration within the same domain 45
- push configuration within the same domain 39
- requirements 37
- within the same domain 37

optimized MSDP duplication

- about 36

optimized synthetic backups

- configuring for deduplication 89
- deduplication 35

OptimizedImage attribute 35

P

pd.conf file

- about 119
- editing 126
- settings 119

pdde-config.log 188

performance

- deduplication 52
- monitoring deduplication rate 137

policies

- changing properties 111–112
- creating 111–112

port usage

- and deduplication 33
- troubleshooting 200

Priority for secondary operations 108

provisioning the deduplication storage 65

PureDisk appliance deduplication 14

PureDisk deduplication 14

PureDisk Deduplication Option

- replacing with media server deduplication 60

Q

queue processing 174

- invoke manually 174

R

rebasing 180

recovery

- deduplication storage server 211

recovery *(continued)*

- from deduplication storage server disk failure 209

Red Hat Linux

- deduplication processes do not start 195

replacing PDDO with NetBackup deduplication 60

replacing the deduplication storage server 207

replication

- between NetBackup domains. *See* Auto Image Replication
- for MSDP 36
- to an alternate NetBackup domain. *See* Auto Image Replication

Replication to remote master. *See* Auto Image Replication

reports

- disk logs 143, 146
- disk pool status 143, 146
- disk storage unit 146

resetting the deduplication registry 132

Resiliency property in Resilient Network host properties 114

Resilient connection

- Resilient Network host properties 113

resilient network connection

- log file 191

Resilient Network host properties 113

- FQDN or IP Address property in 114
- Resiliency property in 114

restores

- at a remote site 183
- how deduplication restores work 59
- specifying the restore server 184

retention periods

- lifecycle and policy-based 110

reverse host name lookup

- prohibiting 193

reverse name lookup 193

S

scaling deduplication 55

scanned field of the job details 141

seeding

- initial 53

server not found error 193

setting deduplication pool attributes 164

spa.cfg file

- parameters for data integrity checking 178

- storage capacity
 - about 68
 - adding 70
 - for deduplicaton 68
 - viewing capacity in container files 143
- storage lifecycle policies
 - Data classification setting 107
 - operations 108
 - Priority for secondary operations 108
 - retention type 110
 - Storage lifecycle policy name 107
 - Suspend secondary operations 108
 - Validate Across Backup Policies button 108
- storage paths
 - about reconfiguring 155
 - changing 155
 - for deduplication 69
- storage requirements
 - for deduplication 66
- storage server
 - about the configuration file 128
 - change the name 155
 - changing properties for deduplication 153
 - changing the name 155
 - components for deduplication 215
 - configuring for deduplication 79
 - deduplication 25
 - define target for Auto Image Replication 49
 - deleting a deduplication 158
 - deleting the deduplication configuration 159
 - determining the state of a deduplication 150
 - editing deduplication configuration file 129
 - getting deduplication configuration 129
 - recovery 211
 - replacing the deduplication host 207
 - setting the deduplication configuration 131
 - viewing 150
 - viewing attributes 151
- storage server configuration
 - getting 129
 - setting 131
- storage server configuration file
 - editing 129
- storage type
 - for deduplication 22
- storage unit
 - configuring for deduplication 83
 - properties for deduplication 85
 - recommendations for deduplication 86

- Storage unit name 85
- Storage unit type 85
- storage units
 - selection in SLP 110
- stream handlers
 - NetBackup 54
- stream rate field of the job details 141
- supported systems 21
- Suspend secondary operations 108

T

- Target retention type 110
- topology of storage 97, 100
- troubleshooting
 - database system error 192
 - deduplication backup jobs fail 195
 - deduplication processes do not start 195
 - general operational problems 198
 - host name lookup 193
 - installation fails on Linux 191
 - no volume appears in disk pool wizard 194
 - server not found error 193

U

- uninstalling media server deduplication 213

V

- Validation Report 108
- viewing deduplication pool attributes 163
- viewing storage server attributes 151
- VM backup 111
- volume manager
 - Veritas Volume Manager for deduplication storage 70

W

- wizards
 - Policy Configuration 111