

# Symantec NetBackup™ Deduplication Guide

UNIX, Windows, Linux

Release 7.6.1



# Symantec NetBackup™ Deduplication Guide

Documentation version: 7.6.1

## Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Symantec as on premises or hosted services. 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 website at the following URL:

[www.symantec.com/business/support/](http://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/](http://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/](http://www.symantec.com/business/support/)

## Customer service

Customer service information is available at the following URL:

[www.symantec.com/business/support/](http://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	<a href="mailto:customercare_apj@symantec.com">customercare_apj@symantec.com</a>
Europe, Middle-East, and Africa	<a href="mailto:semea@symantec.com">semea@symantec.com</a>
North America and Latin America	<a href="mailto:supportsolutions@symantec.com">supportsolutions@symantec.com</a>

# Contents

Technical Support .....	4	
Chapter 1	Introducing the NetBackup media server deduplication option .....	14
	Providing feedback on Beta documentation .....	14
	About the NetBackup deduplication options .....	14
	New MSDP features and enhancements in NetBackup 7.6.1 .....	15
Chapter 2	Planning your deployment .....	17
	Planning your MSDP deployment .....	17
	NetBackup naming conventions .....	19
	About NetBackup media server deduplication .....	19
	About MSDP storage servers .....	21
	About MSDP load balancing servers .....	22
	About MSDP nodes .....	22
	About MSDP server requirements .....	22
	About MSDP unsupported configurations .....	24
	About NetBackup client deduplication .....	24
	About MSDP client deduplication requirements and limitations .....	26
	About MSDP remote office client deduplication .....	26
	About MSDP remote client data security .....	27
	About remote client backup scheduling .....	27
	About the NetBackup Deduplication Engine credentials .....	28
	About the network interface for MSDP .....	28
	About MSDP port usage .....	29
	About MSDP optimized synthetic backups .....	29
	About MSDP and SAN Client .....	30
	About MSDP optimized duplication and replication .....	31
	About MSDP performance .....	31
	How file size may affect the MSDP deduplication rate .....	32
	About MSDP stream handlers .....	32
	MSDP deployment best practices .....	33
	Use fully qualified domain names .....	33
	About scaling MSDP .....	33

	Send initial full backups to the storage server .....	34
	Increase the number of MSDP jobs gradually .....	34
	Introduce MSDP load balancing servers gradually .....	34
	Implement MSDP client deduplication gradually .....	35
	Use MSDP compression and encryption .....	35
	About the optimal number of backup streams for MSDP .....	36
	About storage unit groups for MSDP .....	36
	About protecting the MSDP data .....	37
	Save the MSDP storage server configuration .....	37
	Plan for disk write caching .....	38
Chapter 3	Provisioning the storage .....	39
	About provisioning the storage for MSDP .....	39
	About MSDP storage capacity .....	40
	About MSDP storage and connectivity requirements .....	41
	Fibre Channel and iSCSI comparison for MSDP .....	42
	Do not modify MSDP storage directories and files .....	43
	About adding additional MSDP storage .....	44
	About volume management for NetBackup MSDP .....	44
Chapter 4	Licensing deduplication .....	45
	About licensing MSDP .....	45
	About the MSDP license key .....	45
	Licensing NetBackup MSDP .....	46
Chapter 5	Configuring deduplication .....	47
	Configuring MSDP server-side deduplication .....	49
	Configuring MSDP client-side deduplication .....	51
	About the MSDP Deduplication Multi-Threaded Agent .....	52
	Configuring the Deduplication Multi-Threaded Agent behavior .....	54
	MSDP mtstrm.conf file parameters .....	55
	Configuring deduplication plug-in interaction with the Multi-Threaded Agent .....	59
	About MSDP fingerprinting .....	60
	About the MSDP fingerprint cache .....	60
	Configuring the MSDP fingerprint cache behavior .....	61
	MSDP fingerprint cache behavior options .....	61
	About seeding the MSDP fingerprint cache for remote client deduplication .....	62
	Configuring MSDP fingerprint cache seeding on the client .....	63



Configuring MSDP fingerprint cache seeding on the storage server .....	64
NetBackup seedutil options .....	65
Configuring a storage server for a Media Server Deduplication Pool .....	67
MSDP storage path properties .....	78
MSDP network interface properties .....	80
Configuring a storage server for a PureDisk Deduplication Pool .....	81
About disk pools for NetBackup deduplication .....	91
Configuring a disk pool for deduplication .....	92
Media Server Deduplication Pool properties .....	100
Configuring a Media Server Deduplication Pool storage unit .....	102
Media Server Deduplication Pool storage unit properties .....	103
MSDP storage unit recommendations .....	105
Configuring client attributes for MSDP client-side deduplication .....	106
Disabling MSDP client-side deduplication for a client .....	107
About MSDP compression .....	108
About MSDP encryption .....	110
MSDP compression and encryption settings matrix .....	111
Configuring encryption for MSDP backups .....	112
Configuring encryption for MSDP optimized duplication and replication .....	114
Configuring optimized synthetic backups for MSDP .....	115
About a separate network path for MSDP duplication and replication .....	116
Configuring a separate network path for MSDP duplication and replication .....	117
About MSDP optimized duplication within the same domain .....	118
About the media servers for MSDP optimized duplication within the same domain .....	121
Configuring MSDP optimized duplication within the same NetBackup domain .....	126
Configuring NetBackup optimized duplication or replication behavior .....	130
About MSDP replication to a different domain .....	133
Configuring MSDP replication to a different NetBackup domain .....	135
About NetBackup Auto Image Replication .....	136
About trusted master servers for Auto Image Replication .....	144
Enabling NetBackup clustered master server inter-node authentication .....	145
Configuring a target for MSDP replication to a remote domain .....	146

About configuring MSDP optimized duplication and replication	
bandwidth .....	151
About storage lifecycle policies .....	152
About the storage lifecycle policies required for Auto Image	
Replication .....	153
Creating a storage lifecycle policy .....	155
Storage Lifecycle Policy dialog box settings .....	158
About MSDP backup policy configuration .....	160
Creating a backup policy .....	161
Resilient Network properties .....	162
Resilient connection resource usage .....	165
Specifying resilient connections .....	166
Adding an MSDP load balancing server .....	168
About the MSDP pd.conf configuration file .....	169
Editing the MSDP pd.conf file .....	170
MSDP pd.conf file parameters .....	171
About the MSDP contentrouter.cfg file .....	182
About saving the MSDP storage server configuration .....	183
Saving the MSDP storage server configuration .....	184
Editing an MSDP storage server configuration file .....	184
Setting the MSDP storage server configuration .....	186
About the MSDP host configuration file .....	187
Deleting an MSDP host configuration file .....	187
Resetting the MSDP registry .....	187
About protecting the MSDP catalog .....	188
About the MSDP shadow catalog .....	188
About the MSDP catalog backup policy .....	189
Changing the MSDP shadow catalog path .....	191
Changing the MSDP shadow catalog schedule .....	192
Changing the number of MSDP catalog shadow copies .....	193
Configuring an MSDP catalog backup .....	194
MSDP drcontrol options .....	195
Updating an MSDP catalog backup policy .....	197
 Chapter 6	
Monitoring deduplication activity .....	199
Monitoring the MSDP deduplication rate .....	199
Viewing MSDP job details .....	201
MSDP job details .....	201
About MSDP storage capacity and usage reporting .....	203
About MSDP container files .....	205
Viewing storage usage within MSDP container files .....	205
Viewing MSDP disk reports .....	207

## Chapter 7

About monitoring MSDP processes .....	208
Reporting on Auto Image Replication jobs .....	208
<b>Managing deduplication .....</b>	<b>210</b>
Managing MSDP servers .....	211
Viewing MSDP storage servers .....	211
Determining the MSDP storage server state .....	211
Viewing MSDP storage server attributes .....	212
Setting MSDP storage server attributes .....	213
Changing MSDP storage server properties .....	214
Clearing MSDP storage server attributes .....	215
About changing the MSDP storage server name or storage path .....	216
Changing the MSDP storage server name or storage path .....	216
Removing an MSDP load balancing server .....	218
Deleting an MSDP storage server .....	219
Deleting the MSDP storage server configuration .....	220
About shared memory on Windows MSDP storage servers .....	221
Managing NetBackup Deduplication Engine credentials .....	221
Determining which media servers have deduplication credentials .....	221
Adding NetBackup Deduplication Engine credentials .....	222
Changing NetBackup Deduplication Engine credentials .....	222
Deleting credentials from a load balancing server .....	222
Managing Media Server Deduplication Pools .....	223
Viewing Media Server Deduplication Pools .....	223
Determining the Media Server Deduplication Pool state .....	224
Changing a Media Server Deduplication Pool state .....	224
Viewing Media Server Deduplication Pool attributes .....	224
Setting a Media Server Deduplication Pool attribute .....	225
Changing a Media Server Deduplication Pool properties .....	226
Clearing a Media Server Deduplication Pool attribute .....	231
Determining the MSDP disk volume state .....	232
Changing the MSDP disk volume state .....	232
Inventorying a NetBackup disk pool .....	233
Deleting a Media Server Deduplication Pool .....	234
Deleting backup images .....	234
About MSDP queue processing .....	235
Processing the MSDP transaction queue manually .....	235
About MSDP data integrity checking .....	236
Configuring MSDP data integrity checking behavior .....	237
MSDP data integrity checking configuration parameters .....	239

	About managing MSDP storage read performance .....	241
	About MSDP storage rebasing .....	241
	MSDP server-side rebasing parameters .....	243
	About the MSDP data removal process .....	244
	Resizing the MSDP storage partition .....	245
	How MSDP restores work .....	246
	Configuring MSDP restores directly to a client .....	247
	About restoring files at a remote site .....	247
	About restoring from a backup at a target master domain .....	248
	Specifying the restore server .....	248
Chapter 8	Recovering MSDP .....	250
	About recovering the MSDP catalog .....	250
	Restoring the MSDP catalog from a shadow copy .....	251
	Recovering from an MSDP storage server disk failure .....	253
	Recovering from an MSDP storage server failure .....	254
	Recovering the MSDP storage server after NetBackup catalog recovery .....	257
Chapter 9	Replacing MSDP hosts .....	258
	Replacing the MSDP storage server host computer .....	258
Chapter 10	Uninstalling MSDP .....	261
	About uninstalling MSDP .....	261
	Deactivating MSDP .....	261
Chapter 11	Deduplication architecture .....	263
	MSDP server components .....	263
	Media server deduplication backup process .....	266
	MSDP client components .....	269
	MSDP client-side deduplication backup process .....	269
Chapter 12	Troubleshooting .....	274
	About unified logging .....	274
	About using the vxlogview command to view unified logs .....	275
	Examples of using vxlogview to view unified logs .....	276
	About legacy logging .....	277
	Creating NetBackup log file directories .....	278
	About MSDP log files .....	279
	Troubleshooting MSDP installation issues .....	284

MSDP installation on SUSE Linux fails .....	284
Troubleshooting MSDP configuration issues .....	285
MSDP storage server configuration fails .....	285
MSDP database system error (220) .....	286
MSDP server not found error .....	286
License information failure during MSDP configuration .....	287
The disk pool wizard does not display an MSDP volume .....	288
Troubleshooting MSDP operational issues .....	288
Verify that the MSDP server has sufficient memory .....	289
MSDP backup or duplication job fails .....	289
MSDP client deduplication fails .....	291
MSDP volume state changes to DOWN when volume is unmounted .....	291
MSDP errors, delayed response, hangs .....	292
Cannot delete an MSDP disk pool .....	293
MSDP media open error (83) .....	294
MSDP media write error (84) .....	296
MSDP no images successfully processed (191) .....	299
MSDP storage full conditions .....	299
Troubleshooting MSDP catalog backup .....	300
Viewing MSDP disk errors and events .....	301
MSDP event codes and messages .....	301
Appendix A      Migrating to MSDP storage .....	304
Migrating from PureDisk to the NetBackup MSDP .....	304
Migrating from another storage type to MSDP .....	305
Appendix B      NetBackup appliance deduplication .....	307
About NetBackup appliance deduplication .....	307
About Fibre Channel to a NetBackup 5020 appliance .....	308
Configuring Fibre Channel to a NetBackup 5020 appliance .....	309
Disabling Fibre Channel to a NetBackup 5020 appliance .....	309
Displaying NetBackup 5020 appliance Fibre Channel port information .....	310
Index .....	313

# Introducing the NetBackup media server deduplication option

This chapter includes the following topics:

- [Providing feedback on Beta documentation](#)
- [About the NetBackup deduplication options](#)
- [New MSDP features and enhancements in NetBackup 7.6.1](#)

## Providing feedback on Beta documentation

To provide feedback on this beta documentation, please send us email at the following address:

[DL-ENG-NBU-Early-Release@symantec.com](mailto:DL-ENG-NBU-Early-Release@symantec.com)

## About the NetBackup deduplication options

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 the following benefits:

- Reduce the amount of data that is stored.
- Reduce backup bandwidth.
- Reduce backup windows.
- Reduce infrastructure.

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
Media server deduplication	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 a <b>Media Server Deduplication Pool</b> on the target storage and manages the deduplicated data.  See “ <a href="#">About NetBackup media server deduplication</a> ” on page 19.
Client deduplication	With NetBackup MSDP client deduplication, clients deduplicate their backup data and then send it directly to the storage server, which writes it to the storage. The network traffic is reduced greatly.  See “ <a href="#">About NetBackup client deduplication</a> ” on page 24.

## New MSDP features and enhancements in NetBackup 7.6.1

For the 7.6.1 release, new **Media Server Deduplication Pool** features focus on performance, scalability, and MSDP catalog protection, as follows:

- To improve performance and scalability, NetBackup includes changes to the deduplication database (that is, the *MSDP catalog*). The MSDP database changes in NetBackup 7.6.1 require that the existing database records be converted to a new format during an upgrade installation.  
For information about upgrading and converting your MSDP environment to NetBackup 7.6.1, refer to the *NetBackup Upgrade Guide*:  
<http://www.symantec.com/docs/DOC5332>
- To increase availability, NetBackup introduces a two-tier approach in this release to protect the MSDP catalog, as follows:
  - Daily shadow copies. NetBackup automatically creates copies of the MSDP catalog daily, which provides the first tier of protection. If NetBackup detects corruption in any part the catalog, NetBackup restores that part automatically from the most recent shadow copy.

- **Catalog backup policy.** NetBackup provides a utility that you can use to create a backup policy for the MSDP catalog. It populates the backup policy with appropriate values for your MSDP catalog.

The MSDP catalog backups provide the second tier of catalog protection. The catalog backups are available if the shadow copies are not available or corrupt.

See [“About protecting the MSDP catalog”](#) on page 188.



# Planning your deployment

This chapter includes the following topics:

- [Planning your MSDP deployment](#)
- [NetBackup naming conventions](#)
- [About NetBackup media server deduplication](#)
- [About NetBackup client deduplication](#)
- [About MSDP remote office client deduplication](#)
- [About the NetBackup Deduplication Engine credentials](#)
- [About the network interface for MSDP](#)
- [About MSDP port usage](#)
- [About MSDP optimized synthetic backups](#)
- [About MSDP and SAN Client](#)
- [About MSDP optimized duplication and replication](#)
- [About MSDP performance](#)
- [About MSDP stream handlers](#)
- [MSDP deployment best practices](#)

## Planning your MSDP 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	Determine the storage destination	
Step 2	Determine which type of deduplication to use	See <a href="#">“About NetBackup media server deduplication”</a> on page 19. See <a href="#">“About NetBackup client deduplication”</a> on page 24. See <a href="#">“About MSDP remote office client deduplication”</a> on page 26.
Step 3	Determine the requirements for deduplication hosts	See <a href="#">“About MSDP storage servers”</a> on page 21. See <a href="#">“About MSDP server requirements”</a> on page 22. See <a href="#">“About MSDP client deduplication requirements and limitations”</a> on page 26. See <a href="#">“About the network interface for MSDP”</a> on page 28. See <a href="#">“About MSDP port usage”</a> on page 29. See <a href="#">“About scaling MSDP”</a> on page 33. See <a href="#">“About MSDP performance”</a> on page 31.
Step 4	Determine the credentials for deduplication	See <a href="#">“About the NetBackup Deduplication Engine credentials”</a> on page 28.
Step 5	Read about compression and encryption	See <a href="#">“About MSDP compression”</a> on page 108. See <a href="#">“About MSDP encryption”</a> on page 110.
Step 6	Read about optimized synthetic backups	See <a href="#">“About MSDP optimized synthetic backups”</a> on page 29.
Step 7	Read about deduplication and SAN Client	See <a href="#">“About MSDP and SAN Client”</a> on page 30.
Step 8	Read about optimized duplication and replication	See <a href="#">“About MSDP optimized duplication and replication”</a> on page 31.
Step 9	Read about stream handlers	See <a href="#">“About MSDP stream handlers”</a> on page 32.
Step 10	Read about best practices for implementation	See <a href="#">“MSDP deployment best practices”</a> on page 33.
Step 11	Determine the storage requirements and provision the storage	See <a href="#">“About provisioning the storage for MSDP”</a> on page 39. See <a href="#">“About MSDP storage and connectivity requirements”</a> on page 41. See <a href="#">“About MSDP storage capacity”</a> on page 40. See <a href="#">“MSDP storage path properties”</a> on page 78.

Table 2-1 Deployment overview (*continued*)

Step	Deployment task	Where to find the information
Step 12	License MSDP	See <a href="#">“About licensing MSDP”</a> on page 45.
Step 13	Configure MSDP	See <a href="#">“Configuring MSDP server-side deduplication”</a> on page 49. See <a href="#">“Configuring MSDP client-side deduplication”</a> on page 51.
Step 14	Migrate from other storage to NetBackup deduplication	See <a href="#">“Migrating from another storage type to MSDP”</a> on page 305.

## NetBackup naming conventions

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

- 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 allowed.

---

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

See [“About the NetBackup Deduplication Engine credentials”](#) on page 28.


## About NetBackup media server deduplication

With media server deduplication, the NetBackup client software creates the image of backed up files as for a normal backup. The client sends the backup image to a media server, which hosts the plug-in that duplicates the backup data. The deduplication plug-in breaks the backup image into segments and compares the segments 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**.

[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



More detailed information is available.

See [“About MSDP storage servers”](#) on page 21.

See [“About MSDP load balancing servers”](#) on page 22.

See [“About MSDP nodes”](#) on page 22.

See [“About MSDP server requirements”](#) on page 22.

See [“About MSDP unsupported configurations”](#) on page 24.

See [“MSDP server components”](#) on page 263.

See [“Media server deduplication backup process”](#) on page 266.

## About MSDP storage servers

A storage server is an entity that writes to and reads from the storage. One host functions as the storage server, and only one storage server exists for each NetBackup deduplication node. The host must be a NetBackup media server. Although the storage server components run on a media server, the storage server is a separate logical entity.

See [“About MSDP nodes”](#) on page 22.

The MSDP storage server does the following:

- Receives the backups from clients and then deduplicates the data.
- Receives the deduplicated data from clients or from other media servers.  
 You can configure NetBackup clients and other NetBackup media servers to deduplicate data also. In which case, the storage server only receives the data after it is deduplicated.  
 See [“About NetBackup client deduplication”](#) on page 24.  
 See [“About MSDP load balancing servers”](#) on page 22.
- Writes the deduplicated data to and reads the deduplicated data from the disk storage.
- Manages that storage.
- Manages the deduplication processes.

How many storage servers (and by extension, nodes) you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication, as follows:

- Optimized duplication in the same domain requires at least two deduplication nodes in the same domain. The following are the required storage servers:
  - One for the backup storage, which is the source for the duplication operations.
  - Another to store the copies of the backup images, which are the target for the duplication operations.

See [“About MSDP optimized duplication within the same domain”](#) on page 118.

- Auto Image Replication to another domain requires the following storage servers:

- One for the backups in the originating NetBackup domain. This storage server writes the NetBackup client backups to the storage. It is the source for the duplication operations.
- Another in the remote NetBackup domain for the copies of the backup images. This storage server is the target for the duplication operations that run in the originating domain.

See [“About NetBackup Auto Image Replication”](#) on page 136.

## About MSDP load balancing servers

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.

A NetBackup media server becomes a load balancing server when two things occur:

- 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.

See [“Introduce MSDP load balancing servers gradually”](#) on page 34.

Load balancing servers also perform restore and duplication jobs.

Load balancing servers can be any supported server type for deduplication. They do not have to be the same type as the storage server.

## About MSDP 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 MSDP 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 2-2 shows the minimum requirements for MSDP 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 2-2** MSDP server minimum requirements

Component	Storage server	Load balancing server
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	4 GBs for 4 TBs of storage up to 32 GBs for 64 TBs of storage.	4 GBs.
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 Symantec Support website.</p> <p><a href="http://www.netbackup.com/compatibility">http://www.netbackup.com/compatibility</a></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 following website.</p> <p><a href="http://www.netbackup.com/compatibility">http://www.netbackup.com/compatibility</a></p>

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

---

**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 non-deduplication storage.) If you perform more than 100 backups a day, deduplication operations may affect master server operations.

---

See [“About MSDP performance”](#) on page 31.

See [“About MSDP queue processing”](#) on page 235.

## About MSDP unsupported configurations

The following items describe some configurations that are not supported:

- 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 does not support clustering of deduplication storage servers or load balancing servers.
- Deduplication within each media server deduplication node is supported; global deduplication between nodes is not supported.

## About NetBackup client deduplication

With NetBackup client deduplication (also known as *client-side deduplication*), the client hosts the 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 the segments 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 MSDP remote office client deduplication”](#) on page 26.
- 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



More information is available.

See [“About MSDP client deduplication requirements and limitations”](#) on page 26.

See [“About MSDP remote office client deduplication”](#) on page 26.

See [“MSDP client components”](#) on page 269.

See [“MSDP client-side deduplication backup process”](#) on page 269.

## About MSDP client deduplication requirements and limitations

The clients that deduplicate their own data and the master server must use the same name to resolve the storage server. The name must be the host name under which the NetBackup Deduplication Engine credentials were created. If they do not use the same name, backups fail. In some environments, careful configuration may be required to ensure that the client and the master server use the same name for the storage server. Such environments include those that use VLAN tagging and those that use multi-homed hosts.

NetBackup does not support the following for client-side deduplication:

- Multiple copies per job. For the jobs that specify multiple copies, the backup images are sent to the storage server and may be deduplicated there. Multiple copies are configured in a NetBackup backup policy.
- NDMP hosts. The backup jobs fail if you try to use client-side deduplication for NDMP hosts.

For the systems on which NetBackup supports client-side deduplication, see the NetBackup compatibility lists as the following URL:

<http://www.netbackup.com/compatibility>

The clients that deduplicate their own data conform to the standard NetBackup release level compatibility. The *NetBackup Release Notes* for each release defines the compatibility between NetBackup releases. To take advantage of any new features, improvements, and fixes, Symantec recommends that the clients and the servers be at the same release and revision.

The *NetBackup Release Notes* is available through the following URL:

<http://www.symantec.com/docs/DOC5332>

## About MSDP 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 162.

See [“Specifying resilient connections”](#) on page 166.

The `pd.conf` `FILE_KEEP_ALIVE_INTERVAL` parameter lets you configure the frequency of keep-alive operations on idle sockets.

See [“MSDP pd.conf file parameters”](#) on page 171.

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

See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 62.

## About MSDP remote client 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 MSDP encryption”](#) on page 110.

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 MSDP restores work”](#) on page 246.

## 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 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 the 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:

- The user name and the password can be up to 62 characters in length. The user name and the password cannot be empty or all space characters.
- 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 [)]
  - Less than (<) and greater than (>) sign.
  - Carriage Return (^).
  - Percent sign (%).
  - United States dollar sign (\$).
  - Leading and trailing spaces and quotes.

---

**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 MSDP

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 MSDP storage server name or storage path”](#) on page 216.

---

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

## About MSDP 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 2-3** Deduplication ports

Port	Usage
10082	The NetBackup Deduplication Engine ( <code>spoold</code> ). Open this port between the hosts that deduplicate data. Hosts include load balancing servers and the clients that deduplicate their own data.
10102	The NetBackup Deduplication Manager ( <code>spad</code> ). Open this port between the hosts that deduplicate data. Hosts include load balancing servers and the clients that deduplicate their own data.

## About MSDP optimized synthetic backups

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..

The optimized synthetic backup method provides the following benefits:

- 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.
- Requires no data movement across the network.  
Regular synthetic backups use network traffic.

See [“Configuring optimized synthetic backups for MSDP”](#) on page 115.

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 MSDP storage server attributes”](#) on page 213.

See [“Setting a Media Server Deduplication Pool attribute”](#) on page 225.

**Table 2-4** MSDP requirements and limitations for optimized synthetic backups

What	Description
Requirements	The target storage unit's deduplication pool must be the same deduplication pool on which the source images reside.
Limitations	NetBackup does not support storage unit groups as a destination for optimized synthetic backups. If NetBackup cannot produce the optimized synthetic backup, NetBackup creates the more data-movement intensive synthetic backup.

## About MSDP 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-side 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 MSDP 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 2-5** NetBackup OpenStorage optimized duplication and replication methods

Optimized duplication method	Description
Within the same NetBackup domain	See <a href="#">“About MSDP optimized duplication within the same domain”</a> on page 118.
To a remote NetBackup domain	See <a href="#">“About NetBackup Auto Image Replication”</a> on page 136.

## About MSDP performance

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

[Table 2-6](#) 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 MSDP server requirements”](#) on page 22.

**Table 2-6** MSDP job load performance for an MSDP storage server

When	Description
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"><li>■ The hardware meets minimum requirements. (More capable hardware improves performance.)</li><li>■ No compression. If data is compressed, the CPU usage increases quickly, which reduces the number of concurrent jobs that can be handled.</li><li>■ The deduplication rate is between 50% and 100%. The deduplication rate is the percentage of data already stored so it is not stored again.</li><li>■ The amount of data that is stored is between 30% to 90% of the capacity of the storage.</li></ul>
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"><li>■ The hardware meets minimum requirements. (More capable hardware improves performance.)</li><li>■ The amount of data that is stored is between 85% to 90% of the capacity of the storage.</li></ul> <p>However, the average time to complete the jobs increases significantly.</p>

## How file size may affect the MSDP 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 MSDP 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.



# MSDP deployment best practices

Because Symantec recommends minimum host and network requirements only, deduplication performance may vary greatly depending on your environment. Symantec provides best-practice guidelines to help you use deduplication effectively regardless of the capabilities of your hosts.

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 [“MSDP media write error \(84\)”](#) on page 296.

## About scaling MSDP

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 MSDP:

- 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. Doing so provides information about the total deduplication load. You can then allocate jobs to best balance the load among your hosts.

Deduplication uses the same 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 MSDP fingerprinting”](#) on page 60.

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 MSDP jobs gradually

Symantec recommends that you increase the **Maximum concurrent jobs** value gradually. (The **Maximum concurrent jobs** is a storage unit setting.) Doing so provides information about the total deduplication load. 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. You can then gradually increase the jobs value over time.

See [“About MSDP performance”](#) on page 31.

## Introduce MSDP 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.

See [“About MSDP performance”](#) on page 31.

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 MSDP 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 MSDP 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.  
Doing so provides information about deduplication affects the clients other jobs. It also may be easier to evaluate how your environment handles traffic and easier to troubleshoot any problems

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

## Use MSDP 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 MSDP compression”](#) on page 108.

See [“About MSDP encryption”](#) on page 110.

## About the optimal number of backup streams for MSDP

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 MSDP

You can use a storage unit group as a backup destination for NetBackup MSDP. All of the storage units in the group must have a **Media Server Deduplication Pool** as the storage destination.

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. Symantec recommends that you select the **Failover** method for the **Storage unit selection** type.

**Table 2-7** MSDP requirements and limitations for storage unit groups

What	Description
Requirements	A group must contain storage units of one storage destination type only. That is, a group cannot contain both <b>Media Server Deduplication Pool</b> storage units and storage units with other storage types.

**Table 2-7** MSDP requirements and limitations for storage unit groups  
(continued)

What	Description
Limitations	<p>NetBackup does not support the following for storage unit groups:</p> <ul style="list-style-type: none"><li>■ 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. See <a href="#">“About MSDP optimized duplication within the same domain”</a> on page 118.</li><li>■ Optimized synthetic backups. If NetBackup cannot produce the optimized synthetic backup, NetBackup creates the more data-movement intensive synthetic backup. See <a href="#">“About MSDP optimized synthetic backups”</a> on page 29.</li></ul>

## About protecting the MSDP data

Symantec recommends the following methods to protect the deduplicated backup data:

- Use NetBackup optimized duplication to copy the images to another deduplication node at an 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 [“Configuring MSDP optimized duplication within the same NetBackup domain”](#) on page 126.
- Use NetBackup replication to copy the deduplicated data to another NetBackup domain off-site.  
See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 135.

Symantec also recommends that you back up the MSDP catalog.

See [“About protecting the MSDP catalog”](#) on page 188.

## Save the MSDP 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 MSDP storage server configuration”](#) on page 183.

See [“Saving the MSDP storage server configuration”](#) on page 184.

See [“Editing an MSDP storage server configuration file”](#) on page 184.

## 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.

# Provisioning the storage

This chapter includes the following topics:

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

## About provisioning the storage for MSDP

NetBackup requires that the storage is exposed as a directory path. 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.

See [“About MSDP storage and connectivity requirements”](#) on page 41.

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.

See [“About MSDP optimized duplication within the same domain”](#) on page 118.

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 replication 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 136.

See [“Planning your MSDP deployment”](#) on page 17.

## About MSDP storage capacity

The maximum storage capacity for a single **Media Server Deduplication Pool** 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 and faster disk for the deduplication database (that is, the MSDP catalog). 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 Pool**, you can use more than one media server deduplication node.

See [“About MSDP nodes”](#) on page 22.



# About MSDP storage and connectivity requirements

The following describes the storage and the connectivity 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"> <li>Up to 32 TBs of storage: <ul style="list-style-type: none"> <li>130 MB/sec.</li> <li>200 MB/sec for enterprise-level performance.</li> </ul> </li> <li>32 to 48 TBs of storage: <ul style="list-style-type: none"> <li>200 MB/sec.</li> </ul> <p>Symantec recommends that you store the data and the deduplication database on separate disk, each with 200 MB/sec read or write speed. Neither should be stored on the system disk.</p> </li> <li>48 to 64 TBs of storage: <ul style="list-style-type: none"> <li>250 MB/sec.</li> </ul> <p>Symantec recommends that you store the data and the deduplication database on separate disk, each with 250 MB/sec read or write speed. Neither should be stored on the system disk.</p> </li> </ul> <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"> <li>A dedicated, low latency network for storage with a maximum 0.1-millisecond latency per round trip.</li> <li>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 the Fibre Channel storage networks with at least 4-Gigabit network bandwidth.</li> <li>The storage server should have an HBA or HBAs dedicated to the storage. Those HBAs must have enough bandwidth to satisfy your throughput objectives.</li> </ul> <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> <p>See <a href="#">"Fibre Channel and iSCSI comparison for MSDP"</a> on page 42.</p>

NetBackup requires the 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 **Media Server Deduplication Pool** does not support network attached storage (that is, file based storage protocols), such as CIFS or NFS, for deduplication storage. NetBackup **Media Server Deduplication Pool** does not support the Z file system (ZFS).

The NetBackup compatibility lists are the definitive source for supported operating systems, computers, and peripherals. See the compatibility lists available at the following website:

<http://www.netbackup.com/compatibility>

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

## Fibre Channel and iSCSI comparison for MSDP

Deduplication is a CPU and memory intensive process. It also requires dedicated and high-speed storage connectivity for the best performance. That connectivity helps to ensure the following:

- Consistent storage performance.
- Reduced packet loss during network congestion.
- Reduced storage deadlocks.

The following table compares both the Fibre Channel and the iSCSI characteristics that affect deduplication storage performance. By design, Fibre Channel provides the greatest opportunity to meet performance objectives. To achieve the results that are required for NetBackup MSDP storage, iSCSI may require other optimizations that are described in the following table.

**Table 3-2** Fibre Channel and iSCSI characteristics

Item	Fibre Channel	iSCSI
Genesis	Storage networking architecture that is designed to handle the same block storage format that storage devices use.	Storage network protocol that is built on top of TCP/IP to use the same wiring as the rest of the enterprise.
Protocol	FCP is a thin, single-purpose protocol that provides lossless, in-order frame delivery and low switch latency.	iSCSI is a multiple layer implementation that facilitates data transfers over intranets and long distances. The SCSI protocol expects lossless, in-order delivery, but iSCSI uses TCP/IP, which experiences packet loss and out-of-order delivery.

**Table 3-2** Fibre Channel and iSCSI characteristics (*continued*)

Item	Fibre Channel	iSCSI
Host CPU load	Low. Fibre Channel frame processing is offloaded to dedicated low-latency HBAs.	Higher. Most iSCSI implementations use the host processor to create, send, and interpret storage commands. Therefore, Symantec requires dedicated network interfaces on the storage server to reduce storage server load and reduce latency.
Latency	Low.	Higher.
Flow control	A built-in flow control mechanism that ensures data is sent to a device when it is ready to accept it.	No flow control. However, Ethernet provides a Class of Service (CoS) flow control mechanism. For flow control, Symantec recommends priority-based flow control as defined in the IEEE 802.1Qbb standard.
Deployment	Difficult.	Easy, but more difficult to deploy to meet the criteria for MSDP. The required dedicated network interfaces add to deployment difficult. Other optimizations for carrying storage traffic also add to deployment difficult. Other optimizations include flow control, jumbo framing, and multi-path I/O.

Although Symantec supports iSCSI for connectivity to **Media Server Deduplication Pool** storage, Symantec recommends Fibre Channel. Symantec believes that Fibre Channel provides better performance and stability than iSCSI. iSCSI instability may manifest as status 83 and status 84 error messages.

See [“MSDP media open error \(83\)”](#) on page 294.

See [“MSDP media write error \(84\)”](#) on page 296.

## Do not modify MSDP 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.

- Change the permissions of the directories and files within the deduplication storage directories or database directories.

Failure to follow these directives can result in operational failures and data loss.

## About adding additional MSDP 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 MSDP storage partition](#)” on page 245.

## About volume management for NetBackup MSDP

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

## About licensing MSDP

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, you should determine on which operating systems Symantec supports deduplication. See the NetBackup operating system compatibility list:

<http://www.netbackup.com/compatibility>

## About the MSDP license key

NetBackup deduplication is licensed separately from base NetBackup.

The NetBackup Data Protection Optimization 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 Data Protection Optimization 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 Data Protection Optimization 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 about Accelerator is available.

See the *NetBackup Administrator's Guide, Volume I*:

<http://www.symantec.com/docs/DOC5332>

## Licensing NetBackup MSDP

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 MSDP

- 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 MSDP server-side deduplication](#)
- [Configuring MSDP client-side deduplication](#)
- [About the MSDP Deduplication Multi-Threaded Agent](#)
- [Configuring the Deduplication Multi-Threaded Agent behavior](#)
- [Configuring deduplication plug-in interaction with the Multi-Threaded Agent](#)
- [About MSDP fingerprinting](#)
- [About the MSDP fingerprint cache](#)
- [Configuring the MSDP fingerprint cache behavior](#)
- [About seeding the MSDP fingerprint cache for remote client deduplication](#)
- [Configuring MSDP fingerprint cache seeding on the client](#)
- [Configuring MSDP fingerprint cache seeding on the storage server](#)
- [Configuring a storage server for a Media Server Deduplication Pool](#)
- [Configuring a storage server for a PureDisk Deduplication Pool](#)
- [About disk pools for NetBackup deduplication](#)
- [Configuring a disk pool for deduplication](#)
- [Configuring a Media Server Deduplication Pool storage unit](#)
- [Configuring client attributes for MSDP client-side deduplication](#)
- [Disabling MSDP client-side deduplication for a client](#)

- [About MSDP compression](#)
- [About MSDP encryption](#)
- [MSDP compression and encryption settings matrix](#)
- [Configuring encryption for MSDP backups](#)
- [Configuring encryption for MSDP optimized duplication and replication](#)
- [Configuring optimized synthetic backups for MSDP](#)
- [About a separate network path for MSDP duplication and replication](#)
- [Configuring a separate network path for MSDP duplication and replication](#)
- [About MSDP optimized duplication within the same domain](#)
- [Configuring MSDP optimized duplication within the same NetBackup domain](#)
- [About MSDP replication to a different domain](#)
- [Configuring MSDP replication to a different NetBackup domain](#)
- [About configuring MSDP optimized duplication and replication bandwidth](#)
- [About storage lifecycle policies](#)
- [About the storage lifecycle policies required for Auto Image Replication](#)
- [Creating a storage lifecycle policy](#)
- [About MSDP backup policy configuration](#)
- [Creating a backup policy](#)
- [Resilient Network properties](#)
- [Specifying resilient connections](#)
- [Adding an MSDP load balancing server](#)
- [About the MSDP pd.conf configuration file](#)
- [Editing the MSDP pd.conf file](#)
- [About the MSDP contentrouter.cfg file](#)
- [About saving the MSDP storage server configuration](#)
- [Saving the MSDP storage server configuration](#)
- [Editing an MSDP storage server configuration file](#)



- [Setting the MSDP storage server configuration](#)
- [About the MSDP host configuration file](#)
- [Deleting an MSDP host configuration file](#)
- [Resetting the MSDP registry](#)
- [About protecting the MSDP catalog](#)
- [Changing the MSDP shadow catalog path](#)
- [Changing the MSDP shadow catalog schedule](#)
- [Changing the number of MSDP catalog shadow copies](#)
- [Configuring an MSDP catalog backup](#)
- [Updating an MSDP catalog backup policy](#)

## Configuring MSDP server-side deduplication

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

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

The NetBackup administrator's guide describes how to configure a base NetBackup environment.

See the *NetBackup Administrator's Guide, Volume I*:

<http://www.symantec.com/docs/DOC5332>

**Table 5-1** MSDP configuration tasks

Step	Task	Procedure
Step 1	Install the license key	See <a href="#">"Licensing NetBackup MSDP"</a> on page 46.
Step 2	Create NetBackup log file directories on the master server and the media servers	See <a href="#">"About MSDP log files"</a> on page 279. See <a href="#">"Creating NetBackup log file directories"</a> on page 278.

**Table 5-1** MSDP configuration tasks (*continued*)

Step	Task	Procedure
Step 3	Configure the Deduplication Multi-Threaded Agent behavior	<p>The Deduplication Multi-Threaded Agent uses the default configuration values that control its behavior. You can change those values if you want to do so.</p> <p>See <a href="#">“About the MSDP Deduplication Multi-Threaded Agent”</a> on page 52.</p> <p>See <a href="#">“Configuring the Deduplication Multi-Threaded Agent behavior”</a> on page 54.</p> <p>See <a href="#">“Configuring deduplication plug-in interaction with the Multi-Threaded Agent”</a> on page 59.</p>
Step 4	Configure the fingerprint cache behavior	<p>Configuring the fingerprint cache behavior is optional.</p> <p>See <a href="#">“About the MSDP fingerprint cache”</a> on page 60.</p> <p>See <a href="#">“Configuring the MSDP fingerprint cache behavior”</a> on page 61.</p>
Step 5	Configure a deduplication storage server	<p>How many storage servers you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication. When you configure a storage server, the wizard also lets you configure a disk pool and a storage unit.</p> <p>See <a href="#">“About MSDP storage servers”</a> on page 21.</p> <p>See <a href="#">“MSDP storage path properties”</a> on page 78.</p> <p>See <a href="#">“About MSDP optimized duplication and replication”</a> on page 31.</p> <p>Which type of storage server to configure depends on the storage destination.</p> <p>See <a href="#">“Configuring a storage server for a Media Server Deduplication Pool”</a> on page 67.</p>
Step 6	Configure a disk pool	<p>How many disk pools you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication.</p> <p>If you already configured a disk pool during the storage server configuration, you can skip this step.</p> <p>See <a href="#">“About disk pools for NetBackup deduplication”</a> on page 91.</p> <p>See <a href="#">“Configuring a disk pool for deduplication”</a> on page 92.</p>
Step 7	Configure a storage unit	<p>See <a href="#">“Configuring a Media Server Deduplication Pool storage unit”</a> on page 102.</p>
Step 8	Enable encryption	<p>Encryption is optional.</p> <p>See <a href="#">“Configuring encryption for MSDP backups”</a> on page 112.</p>
Step 9	Configure optimized synthetic backups	<p>Optimized synthetic backups are optional.</p> <p>See <a href="#">“Configuring optimized synthetic backups for MSDP”</a> on page 115.</p>

**Table 5-1** MSDP configuration tasks (*continued*)

Step	Task	Procedure
Step 10	Configure MSDP restore behavior	Optionally, you can configure NetBackup to bypass media servers during restores.  See <a href="#">“How MSDP restores work”</a> on page 246.  See <a href="#">“Configuring MSDP restores directly to a client”</a> on page 247.
Step 11	Configure optimized duplication copy	Optimized duplication is optional.  See <a href="#">“About MSDP optimized duplication within the same domain”</a> on page 118.
Step 12	Configure replication	Replication is optional.  See <a href="#">“About MSDP replication to a different domain”</a> on page 133.
Step 13	Configure a backup policy	Use the deduplication storage unit as the destination for the backup policy. If you configured replication, use the storage lifecycle policy as the storage destination.  See <a href="#">“About MSDP backup policy configuration”</a> on page 160.  See <a href="#">“Creating a backup policy”</a> on page 161.
Step 14	Specify advanced deduplication settings	Advanced settings are optional.  See <a href="#">“About the MSDP pd.conf configuration file”</a> on page 169.  See <a href="#">“Editing the MSDP pd.conf file”</a> on page 170.  See <a href="#">“MSDP pd.conf file parameters”</a> on page 171.
Step 15	Protect the MSDP data and catalog	See <a href="#">“About protecting the MSDP data”</a> on page 37.  See <a href="#">“About protecting the MSDP catalog”</a> on page 188.

## Configuring MSDP client-side deduplication

This topic describes how to configure client deduplication in NetBackup. Media server deduplication must be configured before you can configure client-side deduplication.

See [“Configuring MSDP server-side deduplication”](#) on page 49.

**Table 5-2** Client deduplication configuration tasks

Step	Task	Procedure
Step 1	Configure media server deduplication	See <a href="#">“Configuring MSDP server-side deduplication”</a> on page 49.

**Table 5-2** Client deduplication configuration tasks (*continued*)

Step	Task	Procedure
Step 2	Learn about client deduplication	See <a href="#">“About NetBackup client deduplication”</a> on page 24.
Step 3	Configure a resilient connection for remote office clients	Resilient connections are optional. See <a href="#">“About MSDP remote office client deduplication”</a> on page 26. See <a href="#">“Resilient Network properties”</a> on page 162. See <a href="#">“Specifying resilient connections”</a> on page 166.
Step 4	Enable client-side deduplication	See <a href="#">“Configuring client attributes for MSDP client-side deduplication”</a> on page 106.
Step 5	Configure remote client fingerprint cache seeding	Configuring remote client fingerprint cache seeding is optional. See <a href="#">“Configuring MSDP fingerprint cache seeding on the client”</a> on page 63. See <a href="#">“About seeding the MSDP fingerprint cache for remote client deduplication”</a> on page 62. See <a href="#">“Configuring MSDP fingerprint cache seeding on the storage server”</a> on page 64.
Step 6	Configure client-direct restores	Configuring client-direct restores is optional. If you do not do so, restores travel through the NetBackup media server components. See <a href="#">“Configuring MSDP restores directly to a client”</a> on page 247.

## About the MSDP Deduplication Multi-Threaded Agent

Beginning with the NetBackup 7.6 release, the MSDP deduplication process can use a Multi-Threaded Agent for most data sources. The Multi-Threaded Agent runs alongside the deduplication plug-in on both the clients and the media servers. The agent uses multiple threads for asynchronous network I/O and CPU core calculations. During a backup, this agent receives data from the deduplication plug-in through shared memory and processes it using multiple threads to improve throughput performance. When inactive, the agent uses minimal resources.

The NetBackup Deduplication Multi-Threaded Agent improves backup performance for both client-side deduplication and media server deduplication.

The Deduplication Multi-Threaded Agent uses the default configuration values that control its behavior. You can change those values if you want to do so. The following table describes the Multi-Threaded Agent interactions and behaviors. It also provides links to the topics that describe how to configure those interactions and behaviors.

**Table 5-3** Interactions and behaviors

Interaction	Procedure
Multi-Threaded Agent behavior and resource usage	See <a href="#">“Configuring the Deduplication Multi-Threaded Agent behavior”</a> on page 54.
Whether or not the deduplication plug-in sends backups to the Multi-Threaded Agent	See <a href="#">“Configuring deduplication plug-in interaction with the Multi-Threaded Agent”</a> on page 59.
The clients that should use the Deduplication Multi-Threaded Agent for backups	See <a href="#">“Configuring deduplication plug-in interaction with the Multi-Threaded Agent”</a> on page 59.
The backup policies that should use the Deduplication Multi-Threaded Agent	See <a href="#">“Configuring deduplication plug-in interaction with the Multi-Threaded Agent”</a> on page 59.

[Table 5-4](#) describes the operational notes for MSDP multithreading. If the Multi-Threaded Agent is not used, NetBackup uses the single-threaded mode.

**Table 5-4** Multi-Threaded Agent requirements and limitations

Item	Description
Supported systems	NetBackup supports the Multi-Threaded Agent on Linux, Solaris, and Windows operating systems.
Unsupported use cases	<p>NetBackup does not use the Multi-Threading Agent for the following use cases:</p> <ul style="list-style-type: none"> <li>■ Virtual synthetic backups</li> <li>■ NetBackup Accelerator</li> <li>■ <code>SEGKSIZE</code> is greater than 128 (<code>pd.conf</code> file)</li> <li>■ <code>DONT_SEGMENT_TYPES</code> enabled (<code>pd.conf</code> file)</li> <li>■ <code>MATCH_PDRO</code> = 1 (<code>pd.conf</code> file)</li> </ul> <p>See <a href="#">“MSDP <code>pd.conf</code> file parameters”</a> on page 171.</p>

Table 5-4 Multi-Threaded Agent requirements and limitations (continued)

Item	Description
Policy-based compression or encryption	<p>If NetBackup policy-based compression or encryption is enabled on the backup policy, NetBackup does not use the Deduplication Multi-Threaded Agent.</p> <p>Symantec recommends that you use the MSDP compression and encryption rather than NetBackup policy-based compression and encryption.</p> <p>See <a href="#">“About MSDP compression”</a> on page 108.</p> <p>See <a href="#">“About MSDP encryption”</a> on page 110.</p>

# Configuring the Deduplication Multi-Threaded Agent behavior

The `mtstrm.conf` configuration file controls the behavior of the NetBackup Deduplication Multi-Threaded Agent.

See [“About the MSDP Deduplication Multi-Threaded Agent”](#) on page 52.

If you change the `mtstrm.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 `mtstrm.conf` file on all of the hosts.

## To configure the Multi-Threaded Agent behavior

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

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

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

- 2 To change a behavior, specify a new value.

See [“MSDP `mtstrm.conf` file parameters”](#) on page 55.

- 3 Save and close the file.

- 4 Restart the Multi-Threaded Agent on the host, as follows:

- On UNIX:

```

/usr/opensv/pdde/pdag/bin/mtstrmd -terminate
/usr/opensv/pdde/pdag/bin/mtstrmd

```

- On Windows, use the Windows Services manager. The service name is NetBackup Deduplication Multi-Threaded Agent.

## MSDP mtstrm.conf file parameters

The `mtstrm.conf` configuration file controls the behavior of the Deduplication Multi-threaded Agent. The default values balance performance with resource usage.

A procedure exists that describes how to configure these parameters.

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

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

See [“Configuring the Deduplication Multi-Threaded Agent behavior”](#) on page 54.

The `mtstrm.conf` file is comprised of three sections. The parameters must remain within their sections. For descriptions of the parameters, see the following sections:

- [Logging parameters](#)
- [Process parameters](#)
- [Threads parameters](#)

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

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

### Logging parameters

The following table describes the logging parameters of the `mtstrm.conf` configuration file.

**Table 5-5**      Logging parameters (mtstrm.conf file)

Logging Parameter	Description
LogPath	<p>The directory in which the <code>mtstrmd.log</code> files are created.</p> <p>Default values:</p> <ul style="list-style-type: none"> <li>■ Windows: <code>LogPath=install_path\Veritas\pdde\...\netbackup\logs\pdde</code></li> <li>■ UNIX: <code>LogPath=/var/log/puredisk</code></li> </ul>

**Table 5-5**      Logging parameters (mtstrm.conf file) (*continued*)

Logging Parameter	Description
Logging	<p>Specify what to log:</p> <p>Default value: <code>Logging=short,thread</code>.</p> <p>Possible values:</p> <pre>minimal: Critical, Error, Authentication, Bug short   : all of the above plus Warning long    : all of the above plus Info verbose : all of the above plus Notice full    : all of the above plus Trace messages (everything) none    : disable logging</pre> <p>To enable or disable other logging information, append one of the following to the logging value, without using spaces:</p> <pre>,thread : enable thread ID logging. ,date   : enable date logging. ,timing : enable high-resolution timestamps ,silent : disable logging to console</pre>
Retention	<p>How long to retain log files (in days) before NetBackup deletes them.</p> <p>Default value: <code>Retention=7</code>.</p> <p>Possible values: 0-9, inclusive. Use 0 to keep logs forever.</p>
LogMaxSize	<p>The maximum log size (MB) before NetBackup creates a new log file. The existing log files that are rolled over are renamed <code>mtstrmd.log.&lt;date/time stamp&gt;</code></p> <p>Default value: <code>LogMaxSize=500</code>.</p> <p>Possible value: 1 to the maximum operating system file size in MBs, inclusive.</p>

## Process parameters

The following table describes the process parameters of the `mtstrm.conf` configuration file.



**Table 5-6** Process parameters (mtstrm.conf file)

Process Parameter	Description
MaxConcurrentSessions	<p>The maximum number of concurrent sessions that the Multi-Threaded Agent processes. If it receives a backup job when the <code>MaxConcurrentSessions</code> value is reached, the job runs as a single-threaded job.</p> <p>By default, the deduplication plug-in sends backup jobs to the Multi-Threaded Agent on a first-in, first-out basis. However, you can configure which clients and which backup policies the deduplication plug-in sends to the Multi-Threaded Agent. The <code>MTSTRM_BACKUP_CLIENTS</code> and <code>MTSTRM_BACKUP_POLICIES</code> parameters in the <code>pd.conf</code> control the behavior. Filtering the backup jobs that are sent to the Multi-Threaded Agent can be very helpful on the systems that have many concurrent backup jobs.</p> <p>See <a href="#">“MSDP pd.conf file parameters”</a> on page 171.</p> <p>Default value: <code>MaxConcurrentSessions=</code> (calculated by NetBackup; see the following paragraph).</p> <p>NetBackup configures the value for this parameter during installation or upgrade. The value is the hardware concurrency value of the host divided by the <code>BackupFpThreads</code> value (see <a href="#">Table 5-7</a>). (For the purposes of this parameter, the <i>hardware concurrency</i> is the number of CPUs or cores or hyperthreading units.) On media servers, NetBackup may not use all hardware concurrency for deduplication. Some may be reserved for other server processes.</p> <p>For more information about hardware concurrency, see the <code>pd.conf</code> file <code>MTSTRM_BACKUP_ENABLED</code> parameter description.</p> <p>See <a href="#">“MSDP pd.conf file parameters”</a> on page 171.</p> <p>Possible values: 1-32, inclusive.</p> <p><b>Warning:</b> Symantec recommends that you change this value only after careful consideration of how the change affects your system resources. With default configuration values, each session uses approximately 120 to 150 MBs of memory. The memory that is used is equal to <math>(\text{BackupReadBufferCount} * \text{BackupReadBufferSize}) + (3 * \text{BackupShmBufferSize}) + \text{FpCacheMaxMbSize}</math> (if enabled).</p>
BackupShmBufferSize	<p>The size of the buffers (MB) for shared memory copying. This setting affects three buffers: The shared memory buffer itself, the shared memory receive buffer in the <code>mtstrmd</code> process, and the shared memory send buffer on the client process.</p> <p>Default value: <code>BackupShmBufferSize=2</code> (UNIX) or <code>BackupShmBufferSize=8</code> (Windows).</p> <p>Possible values: 1-16, inclusive.</p>

**Table 5-6** Process parameters (mtstrm.conf file) (*continued*)

Process Parameter	Description
BackupReadBufferSize	<p>The size (MB) of the memory buffer to use per session for read operations from a client during a backup.</p> <p>Default value: BackupReadBufferSize=32.</p> <p>Possible values: 16-128, inclusive.</p>
BackupReadBufferCount	<p>The number of memory buffers to use per session for read operations from a client during a backup.</p> <p>Default value: BackupReadBufferCount=3.</p> <p>Possible values: 1 to 10, inclusive.</p>
BackupBatchSendEnabled	<p>Determines whether to use batch message protocols to send data to the storage server for a backup.</p> <p>Default value: BackupBatchSendEnabled=1.</p> <p>Possible values: 0 (disabled) or 1 (enabled).</p>
FpCacheMaxMbSize	<p>The maximum amount of memory (MB) to use per session for fingerprint caching.</p> <p>Default value: FpCacheMaxMbSize=20.</p> <p>Possible values: 0-1024, inclusive.</p>
SessionCloseTimeout	<p>The amount of time to wait in seconds for threads to finish processing when a session is closed before the agent times-out with an error.</p> <p>Default value: 180.</p> <p>Possible values: 1-3600.</p>
SessionInactiveThreshold	<p>The number of minutes for a session to be idle before NetBackup considers it inactive. NetBackup examines the sessions and closes inactive ones during maintenance operations.</p> <p>Default value: 480.</p> <p>Possible values: 1-1440, inclusive.</p>

## Threads parameters

The following table describes the threads parameters of the `mtstrm.conf` configuration file.

**Table 5-7** Threads parameters (mtstrm.conf file)

Threads Parameter	Description
BackupFpThreads	<p>The number of threads to use per session to fingerprint incoming data.</p> <p>Default value: BackupFpThreads= (calculated by NetBackup; see the following explanation).</p> <p>NetBackup configures the value for this parameter during installation or upgrade. The value is equal to the following hardware concurrency threshold values.</p> <ul style="list-style-type: none"> <li>Windows and Linux: The threshold value is 2.</li> <li>Solaris: The threshold value is 4.</li> </ul> <p>For more information about hardware concurrency, see the <code>pd.conf</code> file <code>MTSTRM_BACKUP_ENABLED</code> parameter description.</p> <p>See “MSDP <code>pd.conf</code> file parameters” on page 171.</p>
BackupSendThreads	<p>The number of threads to use per session to send data to the storage server during a backup operation.</p> <p>Default value: BackupSendThreads=1 for servers and BackupSendThreads=2 for clients.</p> <p>Possible values: 1-32, inclusive.</p>
MaintenanceThreadPeriod	<p>The frequency at which NetBackup performs maintenance operations, in minutes.</p> <p>Default value: 720.</p> <p>Possible values: 0-10080, inclusive. Zero (0) disables maintenance operations.</p>

## Configuring deduplication plug-in interaction with the Multi-Threaded Agent

You can control the interaction between the NetBackup deduplication plug-in and the Multi-Threaded Agent. Several settings in the `pd.conf` file on a host control the interaction. A change in a `pd.conf` file 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 MSDP `pd.conf` configuration file](#)” on page 169.

**To configure deduplication plug-in interaction with the Multi-Threaded Agent**

- 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 change a setting, specify a new value. The following are the settings that control the interaction:
    - `MTSTRM_BACKUP_CLIENTS`
    - `MTSTRM_BACKUP_ENABLED`
    - `MTSTRM_BACKUP_POLICIES`
    - `MTSTRM_IPC_TIMEOUT`
- These settings are defined in another topic.
- See [“MSDP pd.conf file parameters”](#) on page 171.
- 3 Save and close the file.
  - 4 Restart the NetBackup Remote Manager and Monitor Service (`nbrmms`) on the host.

## About MSDP fingerprinting

NetBackup uses a unique identifier to identify each file and each file segment that is backed up. The deduplication plug-in reads the backup image and separates the image into files. The plug-in separates the 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 fingerprint calculations are based on the MD5 algorithm. To prevent MD5 collisions, any segments that have different content but the same MD5 hash key get different fingerprints.

See [“About the MSDP fingerprint cache”](#) on page 60.

See [“Media server deduplication backup process”](#) on page 266.

See [“MSDP client-side deduplication backup process”](#) on page 269.

## About the MSDP fingerprint cache

NetBackup uses *fingerprints* to identify the file segments in the backup data. NetBackup writes only unique data segments to a **Media Server Deduplication Pool**. If a segment already is in storage, NetBackup does not store it again.

See [“About MSDP fingerprinting”](#) on page 60.

The storage server maintains an index cache of the fingerprints in RAM. For each backup job, a client requests a list of the fingerprints from its last backup from the server.

The NetBackup Deduplication Engine (`spoold`) loads a percentage of the fingerprints into the cache at startup. After startup, the Engine loads the remaining fingerprints.

You can configure the cache loading behavior.

See [“Configuring the MSDP fingerprint cache behavior”](#) on page 61.

You can also control the fingerprint cache seeding for clients.

See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 62.

## Configuring the MSDP fingerprint cache behavior

You can configure the cache loading behavior.

See [“About the MSDP fingerprint cache”](#) on page 60.

### To configure MSDP fingerprint cache behavior

- 1 On the storage server, open the `contentrouter.cfg` file in a text editor; it resides in the following directory:
  - (UNIX) `storage_path/etc/puredisk`
  - (Windows) `storage_path\etc\puredisk`
- 2 Edit the parameters that control the behavior.

See [“MSDP fingerprint cache behavior options”](#) on page 61.

## MSDP fingerprint cache behavior options

[Table 5-8](#) describes the parameters that control the behavior. All of these options are in the `contentrouter.cfg` file.

The parameters are stored in the `contentrouter.cfg` file.

See [“About the MSDP contentrouter.cfg file”](#) on page 182.

**Table 5-8** Cache load parameters

Behavior	Description
CacheLoadThreadNum	<p>The number of threads to use to load the remaining fingerprints.</p> <p>The <code>CacheLoadThreadNum</code> in the <code>contentrouter.cfg</code> file controls the number of threads. NetBackup begins loading fingerprints from the next container number after the startup fingerprint loading.</p> <p>The default is one.</p>
MaxCacheSize	<p>The percentage of RAM to use for the fingerprint cache.</p> <p>The <code>MaxCacheSize</code> in the <code>contentrouter.cfg</code> file controls percentage of RAM.</p> <p>The default is 75%.</p>

# About seeding the MSDP fingerprint cache for remote client deduplication

Symantec provides a method for *seeding* the fingerprint cache for a new client. The use case that benefits the most from seeding is the first backup of a remote client over a high latency network such as a WAN. The performance of the first backup is then similar to the performance of an existing client.

An important consideration is the client from which to seed the cache. When you choose a similar client, consider the following:

- If most of the information is the operating system files, use any client with the same operating system.
- If most of the information is data, finding a client with the same data may be unlikely. Therefore, consider physically moving a copy of the data to the datacenter. Back up that data on a similar client, and then use that client and policy for the seed.
- The more similar the clients are, the greater the cache hit rate is.

Two methods exist to configure cache seeding. You can use either method. The following table describes the seeding configuration methods.

**Table 5-9** Seeding configuration methods

Host on which to configure seeding	Description
On the client	Configure seeding on the client for one or only a few clients. See “ <a href="#">Configuring MSDP fingerprint cache seeding on the client</a> ” on page 63.
On the storage server	The use case that benefits the most is many clients to seed, and they can use the fingerprint cache from a single host. See “ <a href="#">Configuring MSDP fingerprint cache seeding on the storage server</a> ” on page 64.

To ensure that NetBackup uses the seeded backup images, the first backup of a client after you configure seeding must be a full backup with a single stream. Specifically, the following two conditions must be met in the backup policy:

- The **Attributes** tab **Allow multiple data streams** attribute must be unchecked.
- The backup selection cannot include any **NEW\_STREAM** directives.

If these two conditions are not met, NetBackup may use multiple streams. If the **Attributes** tab **Limit jobs per policy** is set to a number less than the total number of streams, only those streams use the seeded images to populate the cache. Any streams that are greater than the **Limit jobs per policy** value do not benefit from seeding, and their cache hit rates may be close to 0%.

After the first backup, you can restore the original backup policy parameter settings.

For more information about seeding, see the following Symantec tech note:

<http://www.symantec.com/docs/TECH144437>

See “[About the MSDP fingerprint cache](#)” on page 60.

## Configuring MSDP fingerprint cache seeding on the client

Seeding on the client requires the following:

- A client name
- A policy name
- A date after which to stop using the similar client's fingerprint cache.

Information about when to use this seeding method and how to choose a client from which to seed is available.

See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 62.

---

**Warning:** Do not use this procedure on the storage server or the load balancing server. If you do, it affects all clients that are backed up by that host.

---

### To seed the MSDP fingerprint cache on the client

- ◆ Before the first backup of the remote client, edit the `FP_CACHE_CLIENT_POLICY` parameter in the `pd.conf` file on the remote client.

Specify the setting in the following format:

*clienthostname, backuppolicy, date*

*clienthostname*    The name of the existing similar client from which to seed the cache.

**Note:** NetBackup treats long and short host names differently, so ensure that you use the client name as it appears in the policy that backs it up.

*backuppolicy*      The backup policy for that client.

*date*                The last date in mm/dd/yyyy format to use the fingerprint cache from the existing similar client. After this date, NetBackup uses the fingerprints from the client's own backup.

See [“Editing the MSDP pd.conf file”](#) on page 170.

See [“MSDP pd.conf file parameters”](#) on page 171.

## Configuring MSDP fingerprint cache seeding on the storage server

On the storage server, the NetBackup `seedutil` utility creates a special seeding directory for a client. It populates the seeding directory with image references to another client and policy's backup images. The following is the pathname of the seeding directory:

*database\_path/databases/catalog/2/#pdseed/client\_name*

(By default, NetBackup uses the same path for the storage and the catalog; the *database\_path* and the *storage\_path* are the same. If you configure a separate path for the deduplication database, the paths are different.)



When a backup runs, NetBackup loads the fingerprints from the `#pdseed` directory for the client. (Assuming that no fingerprints exist for that client in the usual catalog location.)

Information about when to use this seeding method and how to choose a client from which to seed is available.

See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 62.

#### To seed the fingerprint cache from the storage server

- 1 Before the first backup of the remote client, specify the clients and the policy in the following format:

UNIX: `/usr/openv/pdde/pdag/bin/seedutil -seed -sclient client_name -spolicy policy_name -dclient destination_client_name`

Windows: `install_path\Veritas\pdde\seedutil -seed -sclient client_name -spolicy policy_name -dclient destination_client_name`

---

**Note:** NetBackup treats long and short host names differently, so ensure that you use the client name as it appears in the policy that backs it up.

---

See [“NetBackup seedutil options”](#) on page 65.

- 2 Repeat the command for each client that you want to seed with fingerprints.
- 3 Verify that the seeding directories for the clients were created by using the following command:

```
seedutil -list_clients
```

- 4 Back up the clients.
- 5 After the client or clients are backed up, remove the seeding directories for the clients. The following is the command syntax:

```
seedutil -clear client_name
```

After one full backup for the client or clients, NetBackup clears the seeding directory automatically. If the first backup fails, the seeded data remains for successive attempts. Although NetBackup clears the seeding directory automatically, Symantec recommends that you clear the client seeding directories manually.

## NetBackup seedutil options

The following is the usage statement for the `seedutil` utility:

```
seedutil [-v log_level] [-seed -sclient source_client_name -spolicy
policy_name -dclient destination_client_name [-backupid backup_id]]
[-clear client_name] [-clear_all] [-list_clients] [-list_images
client_name] [-help]
```

The following items describe the options:

<code>-backupid backup_id</code>	The backup ID from which to copy the data for seeding.
<code>-clear client_name</code>	Clear the contents of the seed directory specified by the <i>client_name</i> .
<code>-clear_all</code>	Clear the contents of all of the seed directories.
<code>-dclient destination_client_name</code>	The name of the new client for which you are seeding the data.
<code>-help</code>	Display help for the command.
<code>-list_clients</code>	List all of the clients that have been configured for seeding.
<code>-list_images client_name</code>	List the contents of the seeding directory for the specified client.
<code>-sclient source_client_name</code>	The client from which to copy the data for seeding.  <b>Note:</b> NetBackup treats long and short host names differently, so ensure that you use the client name as it appears in the policy that backs it up.
<code>-seed</code>	Configure seeding.
<code>-spolicy policy_name</code>	The NetBackup policy that backed up the client that you want to use for the seeding data.
<code>-v log_level</code>	The log level.

The following are the directories in which the command resides:

- **UNIX:** /usr/opensv/pdde/pdag/bin
- **Windows:** C:\Program Files\Veritas\pdde

# Configuring a storage server for a Media Server Deduplication Pool

Configure in this context means to configure a NetBackup media server as a storage server for a **Media Server Deduplication Pool**.

See [“About MSDP storage servers”](#) on page 21.

The type of storage.	Select <b>Media Server Deduplication Pool</b> for the type of disk storage.
The credentials for the deduplication engine.	See <a href="#">“About the NetBackup Deduplication Engine credentials”</a> on page 28.
The storage paths.	See <a href="#">“MSDP storage path properties”</a> on page 78.
The network interface.	See <a href="#">“About the network interface for MSDP”</a> on page 28.
The load balancing servers, if any.	See <a href="#">“About MSDP storage servers”</a> on page 21.

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

To configure a NetBackup storage server for a Media Server Deduplication Pool

- 1 In the **NetBackup Administration Console**, select either **NetBackup Management** or **Media and Device Management**.
- 2 In the right pane, click **Configure Disk Storage Servers**.  
 The **Storage Server Configuration Wizard Welcome** panel appears.

- 3 On the **Welcome** panel, select **Media Server Deduplication Pool** from the drop-down menu.

The types of storage servers that you can configure depend on the options for which you are licensed.

The following is an example of the wizard panel:



After you select **Media Server Deduplication Pool**, click **Next**. The **Add Storage Server** wizard panel appears.


- 4 On the **Add Storage Server** panel, select or enter the appropriate information. The following is an example of the wizard panel:

- Media server** Select the media server that you want to configure as the storage server.
- You can add deduplication load balancing servers on the next wizard panel.
- Username** Enter the user name for the NetBackup Deduplication Engine.
- See [“About the NetBackup Deduplication Engine credentials”](#) on page 28.
- Password** Enter the password for the NetBackup Deduplication Engine.
- Confirm password** To confirm the password, re-enter the password.

After you enter the information, click **Next**.

The **Storage Server Properties** panel appears.

- 5
- On the **Storage Server Properties** panel, enter or select the properties for the deduplication storage server.
- See “MSDP storage path properties” on page 78.
- See “MSDP network interface properties” on page 80.
- The following is an example of the wizard panel:



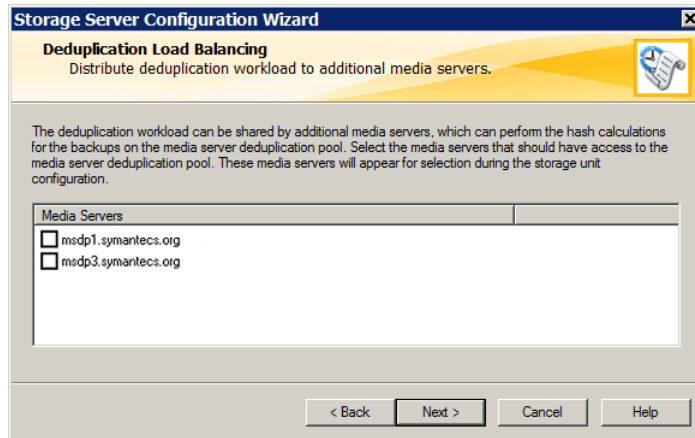
After you click **Next**, the behavior depends on whether you have media servers installed, as follows:

- |                              |   |
|------------------------------|---|
| No media servers installed   | <b>Storage Server Configuration Summary</b> panel appears.<br>Go to step 7.   |
| Media servers are installed. | The <b>Deduplication Load Balancing</b> panel appears.<br>Continue to step 6. |

- 6 On the **Deduplication Load Balancing** panel, select other NetBackup media servers to use for deduplication. Selecting load balancing servers is option.

See [“About MSDP load balancing servers”](#) on page 22.

The following is an example of the wizard panel:

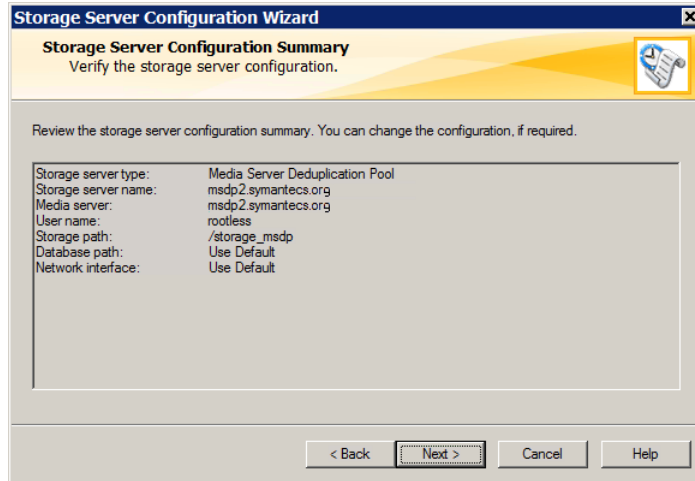


For the media servers that you select, NetBackup enables their deduplication functionality and populates them with the NetBackup Deduplication Engine credentials you entered on a previous panel.

When you click **Next**, the **Storage Server Configuration Summary** panel appears next.

- 7 On the **Storage Server Configuration Summary** panel, verify the selections. If OK, click **Next** to configure the storage server.

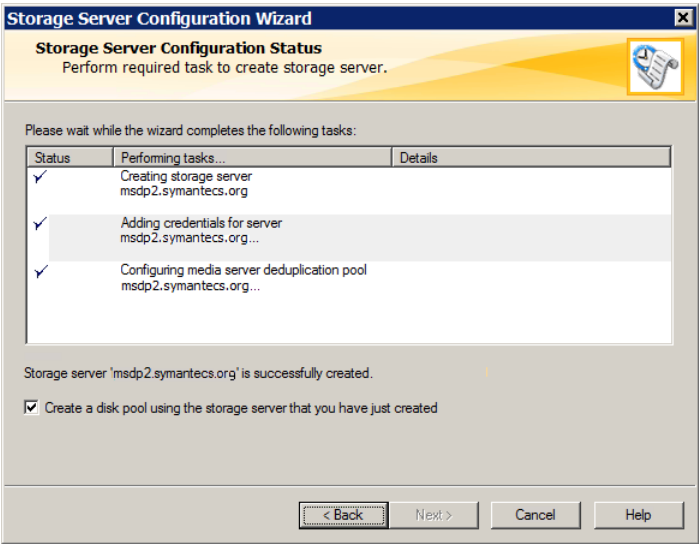
The following is an example of the wizard panel:



The **Storage Server Creation Status** panel appears.



- 8 The **Storage Server Configuration Status** wizard panel describes the status of the operation, as follows:



After the storage server is created, you can do the following:

- Configure a disk pool
- Ensure that **Create a disk pool using the storage server that you have just created** is selected and then click **Next**.  
  
The **Select Volumes** panel appears. Continue to the next step.
- Exit
- Click **Close**.  
  
You can configure a disk pool at another time.

If storage server creation fails, troubleshooting help is available.  
See [“Troubleshooting MSDP configuration issues”](#) on page 285.

- 9 On the **Select Volumes** panel, select the volume for this disk pool. All of storage in the **Storage Path** that you configured in the **Storage Server Configuration Wizard** is exposed as a single volume. The **PureDiskVolume** is a virtual name for that storage.

The following is an example of the wizard panel:

**Storage Server Configuration Wizard**

**Volume Selection**  
Select volumes to use in the disk pool.

Storage server: msdp2.symantecs.org  
Storage server type: PureDisk

Select storage server volumes to add to the disk pool.

Volume Name	Available Space	Raw Size	Replication
<input checked="" type="checkbox"/> PureDiskVolume	723.2441 GB	723.7471 GB	None

Disk pool size:  
 Total available space: 723.2441 GB  
 Total raw size: 723.7471 GB

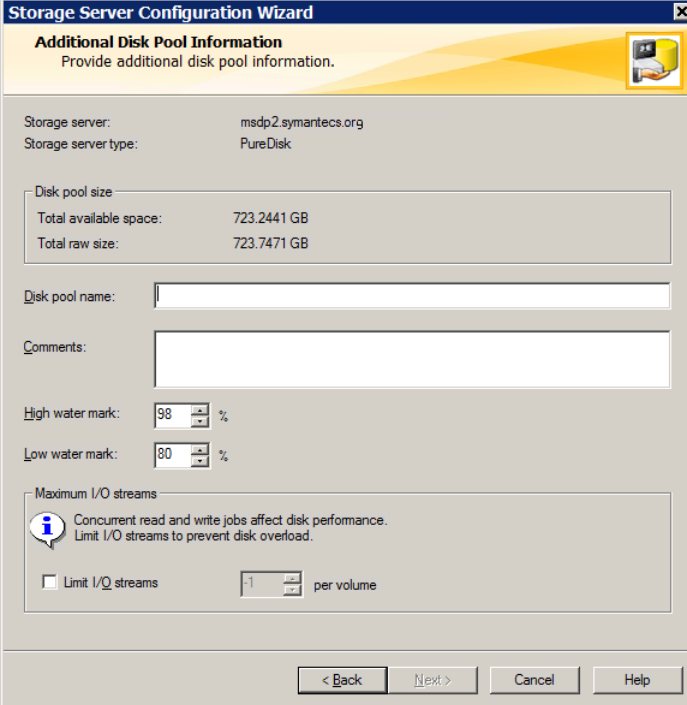
Before selecting a volume, you must validate if it is shared among the storage servers.

After you select the **PureDiskVolume** volume, click **Next**. The **Additional Disk Pool Information** wizard panel appears.

- 10 On the **Additional Disk Pool Information** panel, enter the values for this disk pool.

See “[Media Server Deduplication Pool properties](#)” on page 100.

The following is an example of the wizard panel:



**Storage Server Configuration Wizard**

**Additional Disk Pool Information**  
Provide additional disk pool information.

Storage server: mmdp2.symantecs.org  
Storage server type: PureDisk

— Disk pool size —  
Total available space: 723.2441 GB  
Total raw size: 723.7471 GB

Disk pool name:

Comments:

High water mark: 98 %  
Low water mark: 80 %


— Maximum I/O streams —  
Concurrent read and write jobs affect disk performance.  
Limit I/O streams to prevent disk overload.  
☐ Limit I/O streams -1 per volume

< Back Next > Cancel Help

After you enter the appropriate information or select the necessary options, click **Next**. The **Disk Pool Configuration Summary** wizard panel appears.


- 11 On the **Disk Pool Configuration Summary** panel, verify the selections. If OK, click **Next**.

The following is an example of the wizard panel:



To configure the disk pool, click **Next**. The **Disk Pool Configuration Status** wizard panel appears.

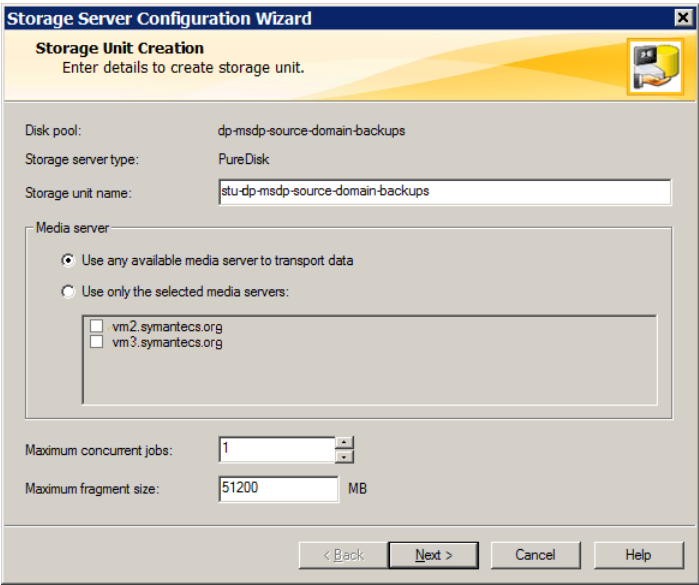
- 12 The **Disk Pool Configuration Status** wizard panel describes the progress of the operation.
- The following is an example of the wizard panel:



After the disk pool is created, you can do the following:

- Configure a storage unit
- Ensure that **Create a storage unit using the disk pool that you have just created** is selected and then click **Next**. The **Storage Unit Creation** wizard panel appears. Continue to the next step.
- Exit
- Click **Close**.
- You can configure one or more storage units later.
- See [“Configuring a Media Server Deduplication Pool storage unit”](#) on page 102.

- 13
- Enter the appropriate information for the storage unit.  
See “Media Server Deduplication Pool storage unit properties” on page 103.  
The following is an example of the wizard panel:



- After you enter the appropriate information or select the necessary options, click **Next** to create the storage unit.
- 14
- After NetBackup configures the storage unit, the **Finished** panel appears. Click **Finish** to exit from the wizard.

## MSDP storage path properties

NetBackup requires that the storage is exposed as a directory path. The following table describes the storage path properties for a **Media Server Deduplication Pool** on the storage server:

Table 5-10MSDP storage path properties

Property	Description
Storage path	<p>The path to the storage. The storage path is the directory in which NetBackup stores the raw backup data. Backup data should not be stored on the system disk.</p> <p>Because the storage requires a directory path, do not use only the root node (/) or drive letter (G:\) as the storage path. (That is, do not mount the storage as a root node (/) or a drive letter (G:\).</p> <p>You can use the following characters in the storage path name:</p> <ul style="list-style-type: none"><li>■ Any of the 26 letters of the International Standards Organization (ISO) Latin-script alphabet, both uppercase (capital) letters and lowercase (small) letters. These are the same letters as the English alphabet.</li><li>■ Any integer from 0 to 9, inclusive.</li><li>■ A space character.</li><li>■ Any of the following characters: UNIX: _ - : . / \ Windows: _ - : . \ (a colon (: ) is allowed only after a drive letter (for example, G:\MSDP_Storage)</li></ul> <p>NetBackup requirements for the deduplication storage paths may affect how you expose the storage.</p> <p>See <a href="#">“About MSDP storage and connectivity requirements”</a> on page 41.</p>
Use alternate path for deduplication database	<p>By default, NetBackup uses the storage path for the MSDP database (that is, the MSDP catalog) location. The MSDP database is different than the NetBackup catalog.</p> <p>Select this option to use a location other than the default for the deduplication database.</p> <p>For performance optimization, Symantec recommends that you use a separate and faster disk for the deduplication database than for the backups.</p>

Table 5-10MSDP storage path properties (continued)

Property	Description
Database path	<p>If you selected <b>Use alternate path for deduplication database</b>, enter the path name for the database. The database should not be stored on the system disk.</p> <p>For performance optimization, Symantec recommends that you use a separate and faster disk for the deduplication database (that is, the MSDP catalog) than for the backup data.</p> <p>You can use the following characters in the path name:</p> <ul style="list-style-type: none"><li>■ Any of the 26 letters of the International Standards Organization (ISO) Latin-script alphabet, both uppercase (capital) letters and lowercase (small) letters. These are the same letters as the English alphabet.</li><li>■ Any integer from 0 to 9, inclusive.</li><li>■ A space character.</li><li>■ Any of the following characters: UNIX: _ - : . / \ Windows: _ - : . \ (a colon (:) is allowed only after a drive letter (for example, C:\MSDP_Storage)</li></ul>

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.

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

MSDP network interface properties

The following table describes the network interface properties for a **Media Server Deduplication Pool** storage server.

**Caution:** You cannot change the network interface after NetBackup configures the deduplication storage server. Therefore, enter the properties carefully.



Table 5-11MSDP network interface properties

Property	Description
Use specific network interface	<p>Select this option to specify a network interface for the backups from and restores to the clients.</p> <p>If you do not specify the network interface, NetBackup uses the operating system host name value.</p> <p>See <a href="#">“About the network interface for MSDP”</a> on page 28.</p>
Interface	<p>If you selected <b>Use specific network interface</b>, enter the interface name.</p>

**Caution:** You cannot change the paths or the network interface after NetBackup configures the deduplication storage server. Therefore, carefully enter the properties

# Configuring a storage server for a PureDisk Deduplication Pool

**Note:** NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

Configure in this context means to configure a PureDisk Storage Pool Authority as a storage server in NetBackup.

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

- The type of storage.  
Select **PureDisk Deduplication Pool** for the type of disk storage.
- The credentials for the PureDisk Storage Pool Authority.
- A media server to use as a backup server and deduplication server. Also, other media servers to use as load balancing servers, if any.

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

To configure a NetBackup storage server for a PureDisk Deduplication Pool by using the wizard

- 1 In the **NetBackup Administration Console**, select either **NetBackup Management** or **Media and Device Management**.

- 2 In the right pane, click **Configure Disk Storage Servers**.

The **Storage Server Configuration Wizard Welcome** panel appears.

- 3 On the **Welcome** panel, select **Media Server Deduplication Pool** from the drop-down menu.

The types of storage servers that you can configure depend on the options for which you are licensed.

The following is an example of the wizard panel:



After you select **Media Server Deduplication Pool**, click **Next**. The **Add Storage Server** wizard panel appears.

- 4 On the **Add Storage Server** panel, select or enter the appropriate information. The following is an example of the wizard panel:

**Media server**

Select a media server to use to query the NetBackup 5000 series appliance for its capabilities. That media server also is used for deduplication.

You can add other deduplication load balancing servers on the next wizard panel.

**Storage server name**

Enter the host name of the NetBackup 5000 series appliance in the **Storage server name** field.

**Username**

Enter the user name for the NetBackup 5000 series appliance.

**Password**

Enter the password for the NetBackup 5000 series appliance.

**Confirm password**

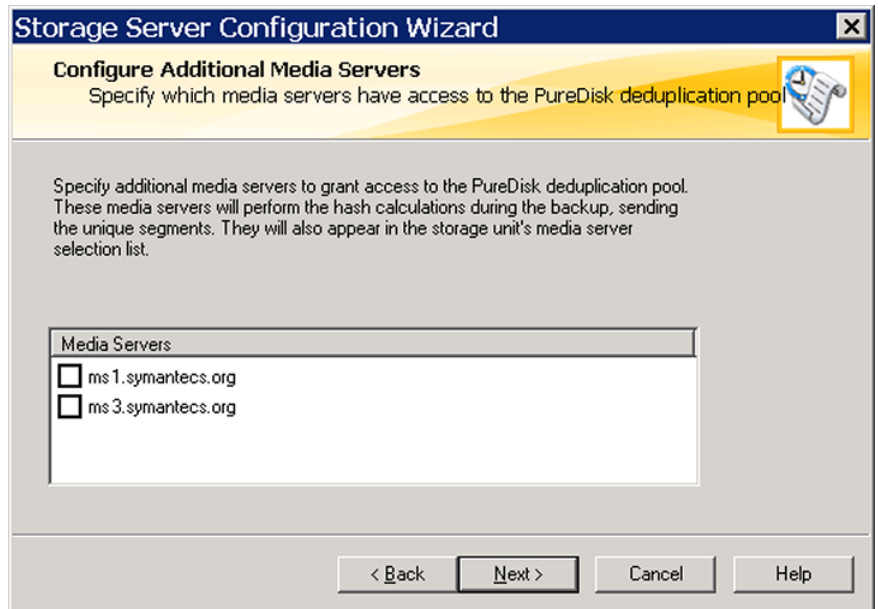
To confirm the password, re-enter the password.

After you enter the information, click **Next**.

The **Configure Additional Media Servers** panel appears.


- 5 On the **Configure Additional Media Servers** panel select other NetBackup media servers to use for deduplication load balancing. The wizard populates the media servers that you select with the Storage Pool Authority credentials you entered on a previous panel. Selecting load balancing servers is option. See [“About MSDP load balancing servers”](#) on page 22.

The following is an example of the wizard panel:



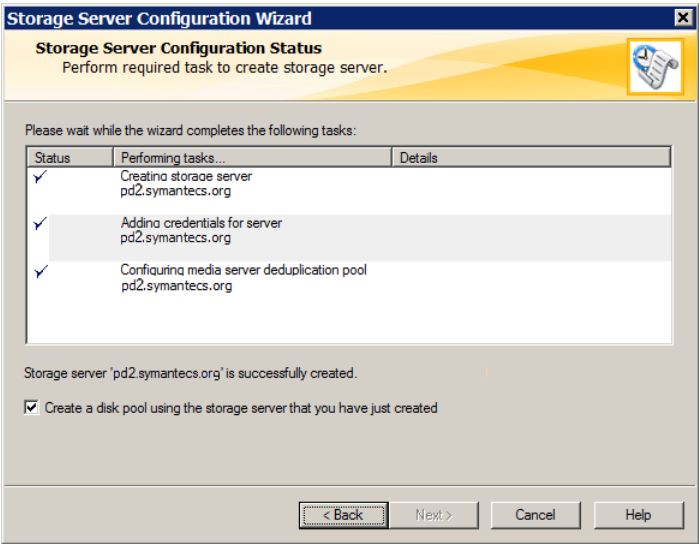
When you click **Next**, the **Storage Server Configuration Summary** panel appears next.

- 6 On the **Storage Server Configuration Summary** panel, verify the selections. If OK, click **Next**.



The **Storage Server Creation Status** panel appears. It describes the status of the wizard actions.

- 7 The **Storage Server Configuration Status** wizard panel describes the status of the operation, as follows:



After the storage server is created, you can do the following:

- Configure a disk pool
- Ensure that **Create a disk pool using the storage server that you have just created** is selected and then click **Next**.  
The **Select Volumes** panel appears. Continue to the next step.
- Exit
- Click **Close**.  
You can configure a disk pool at another time.

If storage server creation fails, troubleshooting help is available.  
See [“Troubleshooting MSDP configuration issues”](#) on page 285.

- 8 On the **Select Volumes** panel, select the volume for this disk pool. All of storage of the NetBackup 5000 series appliance is exposed as a single volume. The **PureDiskVolume** is a virtual name for that storage.

The following is an example of the wizard panel:

**Storage Server Configuration Wizard**

**Volume Selection**  
 Select volumes to use in the disk pool.

Storage server: msdp2.symantecs.org  
 Storage server type: PureDisk

Select storage server volumes to add to the disk pool.

Volume Name	Available Space	Raw Size	Replication
<input checked="" type="checkbox"/> PureDiskVolume	723.2441 GB	723.7471 GB	None

Disk pool size:  
 Total available space: 723.2441 GB  
 Total raw size: 723.7471 GB

Before selecting a volume, you must validate if it is shared among the storage servers.

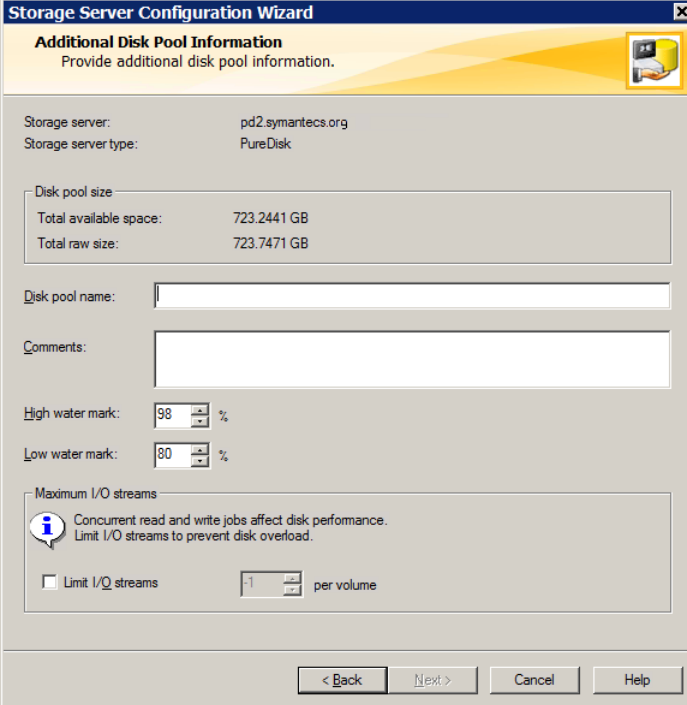
< Back   Next >   Cancel   Help

After you select the **PureDiskVolume** volume, click **Next**. The **Additional Disk Pool Information** wizard panel appears.

- 9 On the **Additional Disk Pool Information** panel, enter the values for this disk pool.

See “[Media Server Deduplication Pool properties](#)” on page 100.

The following is an example of the wizard panel:



**Storage Server Configuration Wizard**

**Additional Disk Pool Information**  
Provide additional disk pool information.

Storage server: pd2.symantecs.org  
Storage server type: PureDisk

Disk pool size  
Total available space: 723.2441 GB  
Total raw size: 723.7471 GB

Disk pool name:

Comments:

High water mark: 98 %  
Low water mark: 80 %

Maximum I/O streams  
Concurrent read and write jobs affect disk performance.  
Limit I/O streams to prevent disk overload.

☐ Limit I/O streams -1 per volume

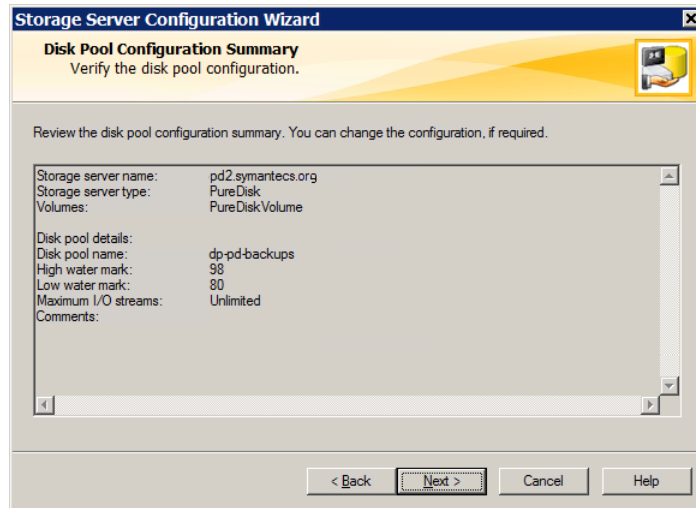
< Back Next > Cancel Help

After you enter the appropriate information or select the necessary options, click **Next**. The **Disk Pool Configuration Summary** wizard panel appears.



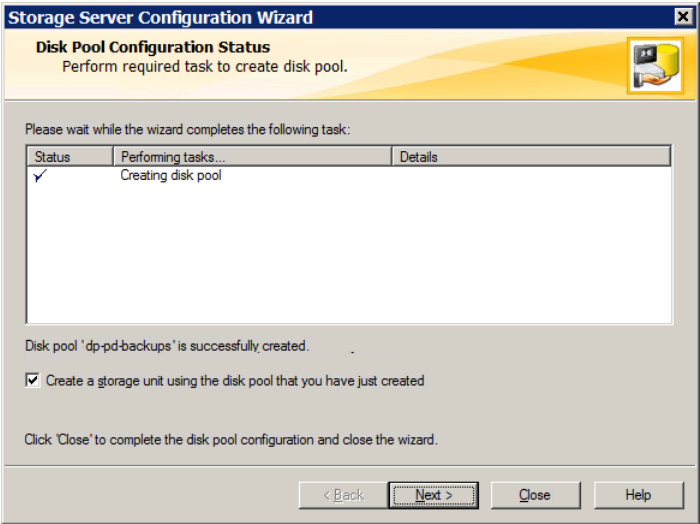
- 10 On the **Disk Pool Configuration Summary** panel, verify the selections. If OK, click **Next**.

The following is an example of the wizard panel:



To configure the disk pool, click **Next**. The **Disk Pool Configuration Status** wizard panel appears.

- 11 The **Disk Pool Configuration Status** wizard panel describes the progress of the operation.
- The following is an example of the wizard panel:



After the disk pool is created, you can do the following:

Configure a storage unit    Ensure that **Create a storage unit using the disk pool that you have just created** is selected and then click **Next**. The **Storage Unit Creation** wizard panel appears. Continue to the next step.

Exit    Click **Close**.

         You can configure one or more storage units later.

         See [“Configuring a Media Server Deduplication Pool storage unit”](#) on page 102.

- 12 Enter the appropriate information for the storage unit.

See “Media Server Deduplication Pool storage unit properties” on page 103.

The following is an example of the wizard panel:

**Storage Server Configuration Wizard**

**Storage Unit Creation**  
Enter details to create storage unit.

Disk pool: dp-pd-backup

Storage server type: PureDisk

Storage unit name: stu-dp-pd-backups

Media server

☒ Use any available media server to transport data

☐ Use only the selected media servers:

- ☐ vm2.symantecs.org
- ☐ vm3.symantecs.org

Maximum concurrent jobs: 1

Maximum fragment size: 51200 MB

< Back Next > Cancel Help

After you enter the appropriate information or select the necessary options, click **Next** to create the storage unit.

- 13 After NetBackup configures the storage unit, the **Finished** panel appears. Click **Finish** to exit from the wizard.

## About disk pools for NetBackup deduplication

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

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.

- A NetBackup **PureDisk Deduplication Pool** represents a PureDisk storage pool on a NetBackup 5000 series appliance. NetBackup deduplicates the data, and PureDisk hosts the storage.

**Note:** NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

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 5-12** 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"><li>■ At least one for the backup storage, which is the source for the duplication operations. The source deduplication pool is in one deduplication node.</li><li>■ 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.</li></ul> <p>See <a href="#">“About MSDP optimized duplication within the same domain”</a> on page 118.</p>
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.</p> <p>See <a href="#">“About the replication topology for Auto Image Replication”</a> on page 141.</p> <p>Auto Image Replication requires the following deduplication pools:</p> <ul style="list-style-type: none"><li>■ 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.</li><li>■ 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.</li></ul> <p>See <a href="#">“About NetBackup Auto Image Replication”</a> on page 136.</p>

## Configuring a disk pool for deduplication

The NetBackup **Storage Server Configuration Wizard** lets you configure one disk pool during storage server configuration. To configure additional disk pools, launch

the **Disk Pool Configuration Wizard**. Before you can configure a NetBackup disk pool, a NetBackup deduplication storage server must exist.

See [“About disk pools for NetBackup deduplication”](#) on page 91.

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

- The type of disk pool: **PureDisk**. **PureDisk** is the type for the following disk pools:
  - A **Media Server Deduplication Pool** on the disk storage that is attached to a NetBackup deduplication media server.
  - A **PureDisk Deduplication Pool**, which represents storage on a NetBackup 5000 series appliance.
- The 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 100.

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

**To configure a deduplication disk pool by using the wizard**


- 1 In the **NetBackup Administration Console**, select either **NetBackup Management** or **Media and Device Management**.
- 2 From the list of wizards in the right pane, click **Configure Disk Pool**.
- 3 Click **Next** on the welcome panel of the wizard.

The **Disk Pool** panel appears.

- 4 On the **Disk Pool** panel, select **PureDisk** in the **Type** windows.

The types of disk pools that you can configure depend on the options for which you are licensed.

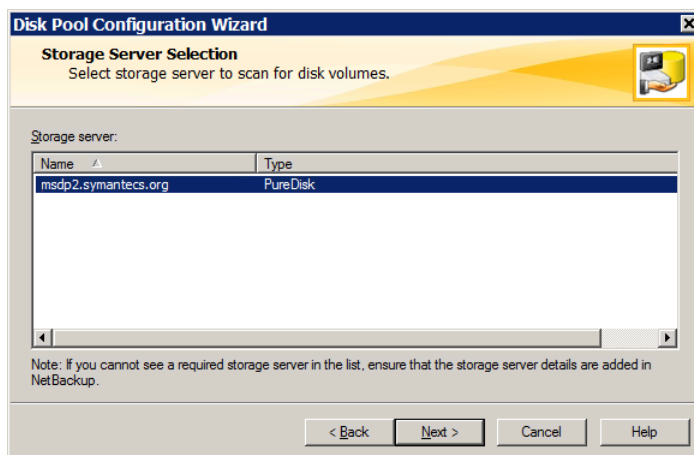
The following is an example of the wizard panel:



After you select the disk pool **Type**, click **Next**. The **Select Storage Server** wizard panel appears.

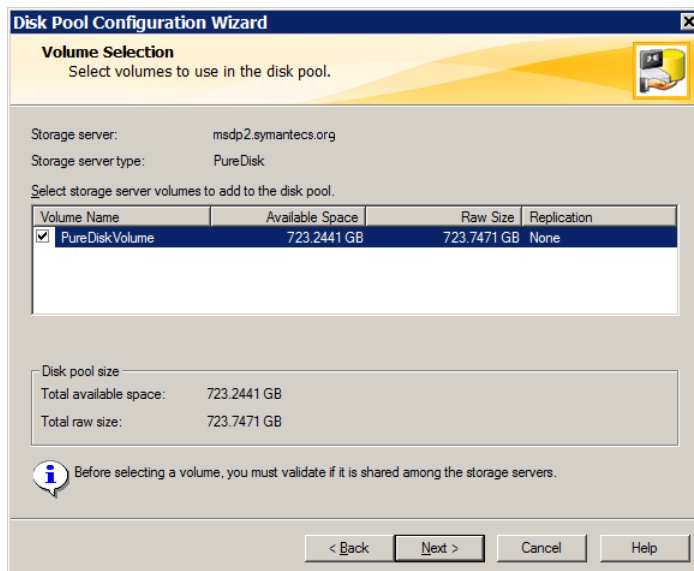
- 5 On the **Select Storage Server** panel, select the storage server for this disk pool. The wizard displays the deduplication storage servers that are configured in your environment. For a **PureDisk Deduplication Pool**, the storage server is the NetBackup 5000 series appliance.

The following is an example of the wizard panel:



Click **Next**. The **Volume Selection** wizard panel appears.

- 6 On the **Volume Selection** panel, select the volume for this disk pool.  
The following is an example of the wizard panel:



**Media Server Deduplication Pool** All of storage in the **Storage Path** that you configured in the **Storage Server Configuration Wizard** is exposed as a single volume. The **PureDiskVolume** is a virtual name for that storage.

**PureDisk Deduplication Pool** All of storage in the NetBackup 5000 series appliance is exposed as a single volume. The **PureDiskVolume** is a virtual name for that storage.

After you select the **PureDiskVolume** volume, click **Next**. The **Additional Disk Pool Information** wizard panel appears.



- 7 On the **Additional Disk Pool Information** panel, enter the values for this disk pool.

See “[Media Server Deduplication Pool properties](#)” on page 100.

The following is an example of the wizard panel:


The screenshot shows a window titled "Disk Pool Configuration Wizard" with a sub-header "Additional Disk Pool Information" and the instruction "Provide additional disk pool information." The window contains the following fields and controls:

- Storage server: mdp2.symantecs.org
- Storage server type: PureDisk
- Disk pool size section:
  - Total available space: 723.2441 GB
  - Total raw size: 723.7471 GB
- Disk pool name: [Empty text box]
- Comments: [Empty text box]
- High water mark: 98 %
- Low water mark: 80 %
- Maximum I/O streams section:
  - Information icon and text: "Concurrent read and write jobs affect disk performance. Limit I/O streams to prevent disk overload."
  - Checkbox: "Limit I/O streams" (unchecked)
  - Value: -1 per volume
- Navigation buttons at the bottom: "< Back", "Next >", "Cancel", and "Help".

After you enter the appropriate information or select the necessary options, click **Next**. The **Disk Pool Configuration Summary** wizard panel appears.

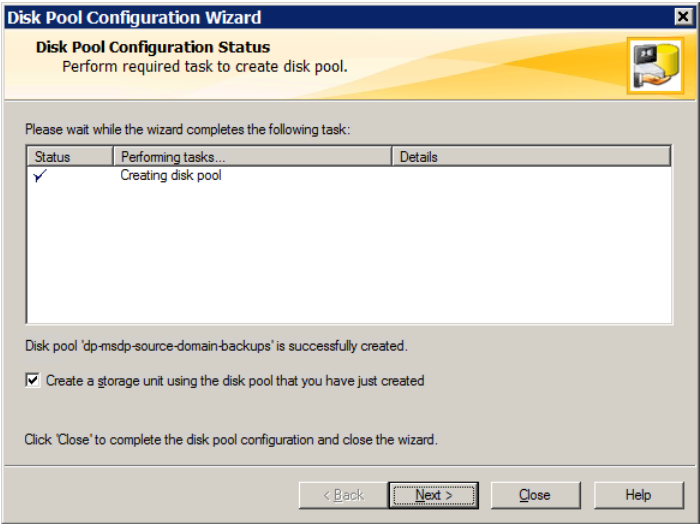
- 8 On the **Disk Pool Configuration Summary** panel, verify the selections. If OK, click **Next**.

The following is an example of the wizard panel:



To configure the disk pool, click **Next**. The **Disk Pool Configuration Status** wizard panel appears.

- 9 The **Disk Pool Configuration Status** wizard panel describes the progress of the operation.
- The following is an example of the wizard panel:



After the disk pool is created, you can do the following:

- Configure a storage unit
- Ensure that **Create a storage unit using the disk pool that you have just created** is selected and then click **Next**. The **Storage Unit Creation** wizard panel appears. Continue to the next step.
- Exit
- Click **Close**.
- You can configure one or more storage units later.
- See [“Configuring a Media Server Deduplication Pool storage unit”](#) on page 102.

- 10 Enter the appropriate information for the storage unit.

See “Media Server Deduplication Pool storage unit properties” on page 103.

The following is an example of the wizard panel:

**Disk Pool Configuration Wizard**  
**Storage Unit Creation**  
Enter details to create storage unit.

Disk pool: dp-msdp-source-domain-backups  
Storage server type: PureDisk  
Storage unit name: stu-dp-msdp-source-domain-backups

Media server

☒ Use any available media server to transport data  
☐ Use only the selected media servers:

☐ vm2.symantecs.org  
☐ vm3.symantecs.org

Maximum concurrent jobs: 1  
Maximum fragment size: 51200 MB

< Back Next > Cancel Help

After you enter the appropriate information or select the necessary options, click **Next** to create the storage unit.

- 11 After NetBackup configures the storage unit, the **Finished** panel appears. Click **Finish** to exit from the wizard.

## Media Server Deduplication Pool properties

Table 5-13 describes the disk pool properties.

**Table 5-13** Media server deduplication pool properties

Property	Description
<b>Storage server</b>	The storage server name. The storage server is the same as the NetBackup media server to which the storage is attached.
<b>Storage server type</b>	For a <b>Media Server Deduplication Pool</b> , the storage type is <b>PureDisk</b> .

Table 5-13 Media server deduplication pool properties (*continued*)

Property	Description
Disk volumes	<p>For a <b>Media Server Deduplication Pool</b>, all disk storage is exposed as a single volume.</p> <p><b>PureDiskVolume</b> is a virtual name for the storage that is contained within the directories you specified for the storage path and the database path.</p>
Total available space	The amount of space available in the disk pool.
Total raw size	The total raw size of the storage in the disk pool.
Disk Pool name	The disk pool name. Enter a name that is unique across your enterprise.
Comments	A comment that is associated with the disk pool.
High water mark	<p>The <b>High water mark</b> indicates that the volume is full. When the volume reaches the <b>High water mark</b>, 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 deduplication pool is full.</p> <p>The <b>High water mark</b> includes the space that is committed to other jobs but not already used.</p> <p>The default is 98%.</p> <p>The default is 98%.</p>
Low water mark	The <b>Low water mark</b> has no affect on the <b>PureDiskVolume</b> .
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>

**Table 5-13** Media server deduplication pool properties (*continued*)

Property	Description
<b>per volume</b>	<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 Media Server Deduplication Pool storage unit

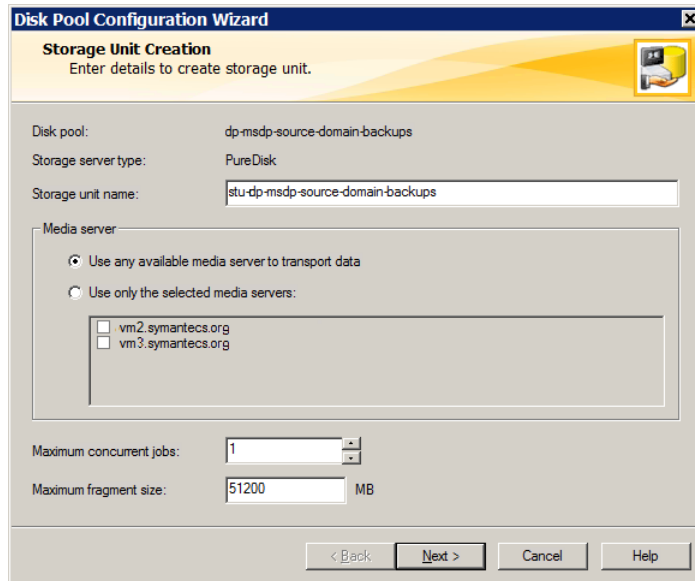
A NetBackup **Media Server Deduplication Pool** storage unit represents the storage in either a **Media Server Deduplication Pool**. Create one or more storage units that reference the disk pool.

See [“About disk pools for NetBackup deduplication”](#) on page 91.

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**.



The screenshot shows the 'Disk Pool Configuration Wizard' window, specifically the 'Storage Unit Creation' step. The window has a yellow header bar with the title 'Storage Unit Creation' and a subtitle 'Enter details to create storage unit.' Below the header, the following fields are visible:

- Disk pool:** dp-msdp-source-domain-backups
- Storage server type:** PureDisk
- Storage unit name:** stu-dp-msdp-source-domain-backups
- Media server:**
  - ☒ Use any available media server to transport data
  - ☐ Use only the selected media servers:
    - ☐ vm2.symantecs.org
    - ☐ vm3.symantecs.org
- Maximum concurrent jobs:** 1
- Maximum fragment size:** 51200 MB

At the bottom of the window, there are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

- 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 [“Media Server Deduplication Pool storage unit properties”](#) on page 103.

## Media Server Deduplication Pool storage unit properties

The following are the configuration options for a storage unit that has a **Media Server Deduplication Pool** as a target.

**Table 5-14** Media Server Deduplication Pool storage unit properties

Property	Description
<b>Storage unit name</b>	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.
<b>Storage unit type</b>	Select <b>Disk</b> as the storage unit type.
<b>Disk type</b>	Select <b>PureDisk</b> for the disk type for a <b>Media Server Deduplication Pool</b> .
<b>Disk pool</b>	Select the disk pool that contains the storage for this storage unit.  All disk pools of the specified <b>Disk type</b> appear in the <b>Disk pool</b> list. If no disk pools are configured, no disk pools appear in the list.
<b>Media server</b>	The <b>Media server</b> setting specifies the NetBackup media servers that can deduplicate the data for this storage unit. Only the deduplication storage server and the load balancing servers appear in the media server list.  Specify the media server or servers as follows: <ul style="list-style-type: none"> <li>■ To allow any server in the media server list to deduplicate data, select <b>Use any available media server</b>.</li> <li>■ To use specific media servers to deduplicate the data, select <b>Only use the following media servers</b>. Then, select the media servers to allow.</li> </ul> NetBackup selects the media server to use when the policy runs.
<b>Maximum fragment size</b>	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.  For more information, see the <i>NetBackup Snapshot Client Administrator's Guide</i> :  <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>



**Table 5-14** Media Server Deduplication Pool storage unit properties (*continued*)

Property	Description
<b>Maximum concurrent jobs</b>	<p>The <b>Maximum concurrent jobs</b> 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 <b>Maximum concurrent jobs</b> 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 <b>Maximum concurrent jobs</b> count.</p> <p><b>Maximum concurrent jobs</b> 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 <b>Maximum concurrent jobs</b>, 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><b>Warning:</b> A Maximum concurrent jobs setting of 0 disables the storage unit.</p>

## MSDP storage unit recommendations

You can use storage unit properties to control how NetBackup performs, as follows:

### Configure a favorable 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 you 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.

# Configuring client attributes for MSDP client-side deduplication

To configure 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 Select the **Client Attributes** properties.
- 5 On the **General** tab of the **Client Attributes** properties, add the clients that you want to deduplicate their own data to the **Clients** list, as follows:
  - Click **Add**.
  - In the **Add Client** dialog box, enter a client name or browse to select a client. Then click **Add**  
Repeat for each client that you want to add.
  - When you finish adding clients, click **Close**.
- 6 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, Volume I*:

<http://www.symantec.com/docs/DOC5332>

See “Disabling MSDP client-side deduplication for a client” on page 107.

## Disabling MSDP 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 MSDP 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 MSDP compression

NetBackup deduplication hosts provide compression for the deduplicated data. It is separate from and different than NetBackup policy-based compression.

Compression is configured by default on all MSDP hosts. Therefore, backups, duplication traffic, and replication traffic are compressed on all MSDP hosts. The data also is compressed on storage.

[Table 5-15](#) describes the compression options.

A different topic describes the interaction of the encryption and the compression settings for MSDP.

See [“MSDP compression and encryption settings matrix”](#) on page 111.

Table 5-15 MSDP compression options

Option	Description
Compression for backups	<p>For backups, the deduplication plug-in compresses the data after it is deduplicated. The data remains compressed during transfer from the plug-in to the NetBackup Deduplication Engine on the storage server. The Deduplication Engine writes the encrypted data to the storage. For restore jobs, the process functions in the reverse direction.</p> <p>The <code>COMPRESSION</code> parameter in the <code>pd.conf</code> file on each MSDP host controls compression and decompression for that host. By default, backup compression is enabled on all MSDP hosts. Therefore, compression and decompression occur on the following hosts as necessary:</p> <ul style="list-style-type: none"> <li>■ The clients that deduplicate their own data (that is, client-side deduplication).</li> <li>■ The load balancing servers.</li> <li>■ The storage server.</li> </ul> <p>MSDP compression cannot occur on normal NetBackup clients (that is, the clients that do not deduplicate their own data).</p> <p><b>Note:</b> Do not enable backup compression by selecting the <b>Compression</b> option on the <b>Attributes</b> tab of the <b>Policy</b> dialog box. If you do, NetBackup compresses the data before it reaches the plug-in that deduplicates it. Consequently, deduplication rates are very low. Also, NetBackup does not use the Deduplication Multi-Threaded Agent if policy-based encryption is configured.</p> <p>See <a href="#">“About the MSDP Deduplication Multi-Threaded Agent”</a> on page 52.</p>
Compression for duplication and replication	<p>For deduplication and replication, the deduplication plug-in compresses the data for transfer. The data remains compressed during transfer from the plug-in to the NetBackup Deduplication Engine on the storage server and remains compressed on the storage.</p> <p>The <code>OPTDUP_COMPRESSION</code> parameter in the <code>pd.conf</code> file controls compression for duplication and replication. By default, duplication and replication compression is enabled on all MSDP hosts. Therefore, duplication and replication compression occurs on the following MSDP servers:</p> <ul style="list-style-type: none"> <li>■ The load balancing servers.</li> <li>■ The storage server.</li> </ul> <p>Duplication and replication compression does not apply to clients.</p> <p>NetBackup chooses the least busy host to initiate and manage each duplication job and replication job. To ensure that compression occurs for all optimized duplication and replication jobs: do not change the default setting of the <code>OPTDUP_COMPRESSION</code> parameter.</p>

See [“About the MSDP pd.conf configuration file”](#) on page 169.

See [“Use MSDP compression and encryption”](#) on page 35.

# About MSDP encryption

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

[Table 5-16](#) describes the encryption options.

A different topic describes the interaction of the encryption and the compression settings for MSDP.

See [“MSDP compression and encryption settings matrix”](#) on page 111.

**Table 5-16** MSDP encryption options

Option	Description
Backup encryption	<p>For backups, the deduplication plug-in encrypts the data after it is deduplicated. The data remains encrypted during transfer from the plug-in to the NetBackup Deduplication Engine on the storage server. The Deduplication Engine writes the encrypted data to the storage. For restore jobs, the process functions in the reverse direction.</p> <p>The MSDP <code>pd.conf</code> file <code>ENCRYPTION</code> parameter controls backup encryption for individual hosts. By default, backup encryption is disabled on all MSDP hosts. If you want backup encryption, you must enable it on the following MSDP hosts:</p> <ul style="list-style-type: none"><li>■ The clients that deduplicate their own data (that is, client-side deduplication).</li><li>■ The MSDP load balancing servers.</li><li>■ The MSDP storage server.</li></ul> <p>See <a href="#">“Configuring encryption for MSDP backups”</a> on page 112.</p> <p><b>Note:</b> Do not enable backup encryption by selecting the <b>Encryption</b> option on the <b>Attributes</b> tab of the <b>Policy</b> dialog box. If you do, NetBackup encrypts the data before it reaches the plug-in that deduplicates it. Consequently, deduplication rates are very low. Also, NetBackup does not use the Deduplication Multi-Threaded Agent if policy-based encryption is configured.</p> <p>See <a href="#">“About the MSDP Deduplication Multi-Threaded Agent”</a> on page 52.</p>

Table 5-16 MSDP encryption options (*continued*)

Option	Description
Duplication and replication encryption	<p>For deduplication and replication, the deduplication plug-in on MSDP servers encrypts the data for transfer. The data is encrypted during transfer from the plug-in to the NetBackup Deduplication Engine on the target storage server and remains encrypted on the target storage.</p> <p>The MSDP <code>pd.conf</code> file <code>OPTDUP_ENCRYPTION</code> parameter controls duplication and replication encryption for individual hosts. By default, duplication and replication encryption is disabled on the MSDP storage server and on the MSDP load balancing servers. If you want duplication and replication encryption, you must enable it on the following MSDP servers:</p> <ul style="list-style-type: none"> <li>■ The load balancing servers.</li> <li>■ The storage server.</li> </ul> <p>Duplication and replication encryption does not apply to clients.</p> <p>NetBackup chooses the least busy host to initiate and manage each duplication job and replication job.</p> <p>See <a href="#">“Configuring encryption for MSDP optimized duplication and replication”</a> on page 114.</p>

## MSDP compression and encryption settings matrix

Four MSDP `pd.conf` file parameters control the compression and the encryption for individual hosts. Table 5-17 shows the matrix for the outcomes for the compression and the encryption parameters.

Table 5-17 Compression and encryption settings and outcomes

Parameters	Default: Compress both the backups and duplication and replication	Settings A: Compress and encrypt the backups	Settings B: Compress the backups and compress and encrypt duplication and replication	Settings C: Compress and encrypt backups and compress duplication and replication	Settings D: Compress and encrypt both backups and duplication and replication
ENCRYPTION	0	1	0	1	1
COMPRESSION	1	1	1	1	1
OPTDUP_ENCRYPTION	0	0	1	0	1
OPTDUP_COMPRESSION	1	0	1	1	1

**Table 5-17** Compression and encryption settings and outcomes (*continued*)

Parameters	Default: Compress both the backups and duplication and replication	Settings A: Compress and encrypt the backups	Settings B: Compress the backups and compress and encrypt duplication and replication	Settings C: Compress and encrypt backups and compress duplication and replication	Settings D: Compress and encrypt both backups and duplication and replication
Outcomes on the disk pools	Backup MSDP: Encryption: 0 Compression: 1  Target MSDP: Encryption: 0 Compression: 1	Backup MSDP: Encryption: 1 Compression: 1  Target MSDP: Encryption: 1 Compression: 1	Backup MSDP: Encryption: 0 Compression: 1  Target MSDP: Encryption: 1 Compression: 1	Backup MSDP: Encryption: 1 Compression: 1  Target MSDP: Encryption: 1 Compression: 1	Backup MSDP: Encryption: 1 Compression: 1  Target MSDP: Encryption: 1 Compression: 1
Notes		See the following note.		See the following note.	See the following note.

---

**Note:** Settings A and C have the same outcome on target storage as setting D because of the interaction of compression and encryption for the backups. If the backups are encrypted and compressed, they are also encrypted and compressed during optimized duplication and replication and encrypted and compressed on target storage. The `OPTDUP_ENCRYPTION` setting is ignored.

---

For client-side deduplication clients, a parameter on the storage server can override the `pd.conf` file `ENCRYPTION` parameter.

See [“Configuring encryption for MSDP backups”](#) on page 112.

See [“About MSDP compression”](#) on page 108.

See [“About MSDP encryption”](#) on page 110.

See [“About the MSDP `pd.conf` configuration file”](#) on page 169.

## Configuring encryption for MSDP backups

Two procedures exist to configure encryption during backups for MSDP, as follows:



- |  |  |
|--|--|
| Configure encryption on individual hosts           | <p>Use this procedure to configure encryption on individual MSDP hosts.</p> <p>The <code>ENCRYPTION</code> parameter in the MSDP <code>pd.conf</code> file controls encryption for that host. The parameter applies only to the host on which you modify the <code>pd.conf</code>, as follows:</p> <p>See <a href="#">“To configure backup encryption on a single host”</a> on page 113.</p>   |
| Configure encryption for all Client Direct clients | <p>Use this procedure to configure encryption for all of your clients that deduplicate their own data (that is, client-side deduplication). If you use this procedure, you do not have to configure each client-side deduplication client individually.</p> <p>The <code>ServerOptions</code> parameter in the MSDP <code>contentrouter.cfg</code> file controls encryption for all client-side deduplication clients. This parameter supersedes the <code>pd.conf</code> file <code>ENCRYPTION</code> setting on client-side deduplication hosts.</p> <p>See <a href="#">“To configure backup encryption on all client-side deduplication clients”</a> on page 114.</p> |

To ensure that encryption occurs for all backups jobs, configure it on all MSDP hosts. MSDP hosts include the MSDP storage server, the MSDP load balancing servers, and the NetBackup Client Direct deduplication clients.

See [“About MSDP encryption”](#) on page 110.

#### To configure backup 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 [“MSDP pd.conf file parameters”](#) on page 171.

- 2 For the line that begins with `#ENCRYPTION`, remove the pound sign ( or hash sign, #) in column 1.
- 3 In that same 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 On the client-side deduplication clients and on the MSDP load balancing servers, ensure that the `LOCAL_SETTINGS` parameter in the `pd.conf` file is set to 1. Doing so ensures that the setting on the current host has precedence over the server setting.
- 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.

#### To configure backup encryption on all client-side deduplication clients

- 1 On the storage server, open the `contentrouter.cfg` file in a text editor; it resides in the following directory:
  - (UNIX) `storage_path/etc/puredisk`
  - (Windows) `storage_path\etc\puredisk`
- 2 Add `agent_crypt` to the `ServerOptions` line of the file. The following line is an example:

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

## Configuring encryption for MSDP optimized duplication and replication

The `OPTDUP_ENCRYPTION` parameter in the `pd.conf` file on the MSDP host controls duplication and replication encryption for that host. The data that is encrypted during transfer remains encrypted on the target storage.

Use this procedure to configure encryption for optimized duplication and replication on MSDP storage servers and on MSDP load balancing servers. To ensure that encryption occurs for all optimized duplication and replication jobs, configure it on all MSDP servers.

By default, optimized duplication encryption is disabled on all MSDP hosts.

See [“About MSDP encryption”](#) on page 110.

#### To configure replication and duplication encryption on an MSDP server

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

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

- (UNIX) `/usr/opensv/lib/ost-plugins/`

- (Windows) `install_path\Veritas\NetBackup\bin\ost-plugins`
- 2 For the line that begins with `#OPTDUP_ENCRYPTION`, remove the pound sign (or hash sign, `#`) in column 1.
  - 3 In that same 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.

See [“MSDP pd.conf file parameters”](#) on page 171.

- 4 On load balancing servers, ensure that the `LOCAL_SETTINGS` parameter is set to `1`. Doing so ensures that the `ENCRYPTION` setting on the load balancing server is used.
- 5 Save and close the file.
- 6 Restart the NetBackup Remote Manager and Monitor Service (`nbrmms`) on the host.

## Configuring optimized synthetic backups for MSDP

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

See [“About MSDP optimized synthetic backups”](#) on page 29.

**Table 5-18** To configure optimized synthetic backups for MSDP

Step	Task	Procedure
Step 1	If your storage server is at NetBackup release 7.0, set the <b>OptimizedImage</b> attribute on the storage server.	The <b>OptimizedImage</b> attribute is set by default beginning with the NetBackup 7.1 release.  See <a href="#">“Setting MSDP storage server attributes”</a> on page 213.
Step 2	If your storage server is at NetBackup release 7.0, set the <b>OptimizedImage</b> attribute on your existing deduplication pools. (Any deduplication pools that you create after you set the storage server attribute inherit the new functionality.)	The <b>OptimizedImage</b> attribute is set by default beginning with the NetBackup 7.1 release.  See <a href="#">“Setting a Media Server Deduplication Pool attribute”</a> on page 225.

**Table 5-18** To configure optimized synthetic backups for MSDP (*continued*)

Step	Task	Procedure
Step 3	Configure a <b>Standard</b> or <b>MS-Windows</b> backup policy. Select the <b>Synthetic backup</b> attribute on the <b>Schedule Attribute</b> tab of the backup policy.	See <a href="#">“Creating a backup policy”</a> on page 161.

See [“Setting MSDP storage server attributes”](#) on page 213.

See [“Creating a backup policy”](#) on page 161.

# About a separate network path for MSDP duplication and replication

You can use a different network for MSDP duplication and replication traffic rather than the one you use for MSDP backups. Both the duplication and the replication data traffic and the control traffic travel over the separate network. Your MSDP traffic uses two different networks, as follows:

Backups and restores	<p>For backups and restores, NetBackup uses the network interface that was configured during the storage server configuration.</p> <p>Both the backup and restore traffic and the control traffic travel over the <i>backup</i> network.</p> <p>See <a href="#">“About the network interface for MSDP”</a> on page 28.</p>
Duplication and replication	<p>For the duplication and the replication traffic, configure your host operating systems to use a different network than the one you use for backups and restores.</p> <p>Both the duplication and the replication data traffic and the control traffic travel over the <i>duplication and replication</i> network.</p> <p>See <a href="#">“Configuring a separate network path for MSDP duplication and replication”</a> on page 117.</p> <p>When you configure the optimized duplication or replication target, ensure that you select the host name that represents the <i>duplication and replication</i> network.</p>

See [“About MSDP optimized duplication within the same domain”](#) on page 118.

See [“About MSDP replication to a different domain”](#) on page 133.

# Configuring a separate network path for MSDP duplication and replication

You can use a different network for MSDP duplication and replication traffic rather than the one you use for MSDP backups. Both the duplication and the replication data traffic and the control traffic travel over the separate network.

See [“About a separate network path for MSDP duplication and replication”](#) on page 116.

This procedure describes how to use the storage servers `hosts` files to route the traffic onto the separate network.

The following are the prerequisites:

- Both the source and the destination storage servers must have a network interface card that is dedicated to the other network.
- The separate network must be operational and using the dedicated network interface cards on the source and the destination storage servers.
- On UNIX MSDP storage servers, ensure that the Name Service Switch first checks the local `hosts` file for before querying the Domain Name System (DNS). See the operating system documentation for information about the Name Service Switch.

## To configure a separate network path for MSDP duplication and replication

- 1 On the source storage server, add the destination storage servers's dedicated network interface to the operating system `hosts` file. If *TargetStorageServer* is the name of the destination host on the network that is dedicated for duplication, the following is an example of the `hosts` entry in IPv4 notation:

```
10.10.10.1 TargetStorageServer. TargetStorageServer
```

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

- 2 On the destination storage server, add the source storage servers's dedicated network interface to the operating system `hosts` file. If *SourceStorageServer* is the name of the source host on the network that is dedicated for duplication, the following is an example of the `hosts` entry in IPv4 notation:

```
10.80.25.66 SourceStorageServer. SourceStorageServer
```

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

- 3 To force the changes to take effect immediately, flush the DNS cache. See the operating system documentation for how to flush the DNS cache.
- 4 From each host, use the `ping` command to verify that each host resolves the name of the other host.

```
SourceStorageServer.> ping TargetStorageServer.  
TargetStorageServer.> ping SourceStorageServer.
```

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

- 5 When you configure the target storage server, ensure that you select the host name that represents the alternate network path.

# About MSDP optimized duplication within the same domain

Optimized duplication within the same domain copies the deduplicated backup images between **Media Server Deduplication Pools** and **PureDisk Deduplication Pools** within the same domain. The source and the destination storage must use the same NetBackup master server.

**Note:** A **PureDisk Deduplication Pool** in NetBackup represents the storage on a NetBackup 5000 series appliance. NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

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 table describes the supported sources and targets.

Table 5-19 MSDP optimized duplication source and targets

Source storage	Target storage
Media Server Deduplication Pool	Media Server Deduplication Pool
Media Server Deduplication Pool	PureDisk Deduplication Pool
PureDisk Deduplication Pool	PureDisk Deduplication Pool

**Table 5-19** MSDP optimized duplication source and targets (*continued*)

Source storage	Target storage
PureDisk Deduplication Pool	Media Server Deduplication Pool

By default, NetBackup does not retry failed optimized duplication jobs invoked by NetBackup Vault of the `bpduplicate` command. You can change that behavior.

See “[Configuring NetBackup optimized duplication or replication behavior](#)” on page 130.

You can use a separate network for the duplication traffic.

See “[About a separate network path for MSDP duplication and replication](#)” on page 116.

See “[Configuring MSDP optimized duplication within the same NetBackup domain](#)” on page 126.

Review the following requirements and limitations.

## About MSDP optimized duplication requirements

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

- The source storage and the destination storage must have at least one media server in common.  
See “[About the media servers for MSDP optimized duplication within the same domain](#)” on page 121.
- 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 in common, the optimized duplication job fails.  
For more information about media server load balancing, see the *NetBackup Administrator's Guide, Volume I*:  
<http://www.symantec.com/docs/DOC5332>
- The destination storage unit cannot be the same as the source storage unit.

## About MSDP optimized duplication limitations

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

- 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 NetBackup optimized duplication or replication behavior”](#) on page 130.

- NetBackup does not support MSDP optimized duplication to 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**.
- For optimized duplication between two **PureDisk Deduplication Pools**, NetBackup reports that the deduplication rate is 100%. However, deduplication does not occur during optimized duplication. Therefore, ignore the deduplication rate statistics.

---

**Note:** NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

---

- The copy operation uses the maximum fragment size of the source storage unit, not the setting for the destination storage unit. The optimized duplication copies the image fragments as is. For greater efficiency, the duplication does not resize and reshuffle the images into a different set of fragments on the destination storage unit.
- If the destination is a **PureDisk Deduplication Pool**, that PureDisk storage domain can be the destination of duplication from one NetBackup domain only.



If you try to duplicate from MSDP storage in more than one NetBackup domain, optimized duplication does not function correctly.

---

**Note:** NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

---

## About the media servers for MSDP optimized 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 duplication 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, as follows:

- 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 5-1](#) shows a push configuration for optimized duplication 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 5-1 Push duplication environment




Figure 5-2 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 5-2** Storage unit settings for backups to MSDP\_L

The screenshot shows a 'New Storage Unit' dialog box with the following settings:

- Storage unit name:** STU-Backups
- Storage unit type:** Disk (selected from a dropdown menu)
- On demand only:** ☒ (checked)
- Disk type:** PureDisk (selected from a dropdown menu)
- Properties section:**
  - Disk Pool:** MSDP\_L (selected from a dropdown menu, with a 'Properties' button next to it)
  - Media Server:**
    - ☒ Use any available media server
    - ☐ Only use the following media servers
  - Media Servers list:**
    - ☐ StorageServer-L
    - ☐ LB\_L1
    - ☐ LB\_L2

Figure 5-3 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 5-3** 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 **Media Server Deduplication Pool**. In such a case, data travels across your WAN.

You can use a load balancing server when you duplicate between two NetBackup deduplication pools.

## About MSDP pull duplication within the same domain

Figure 5-4 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 5-4 Pull duplication

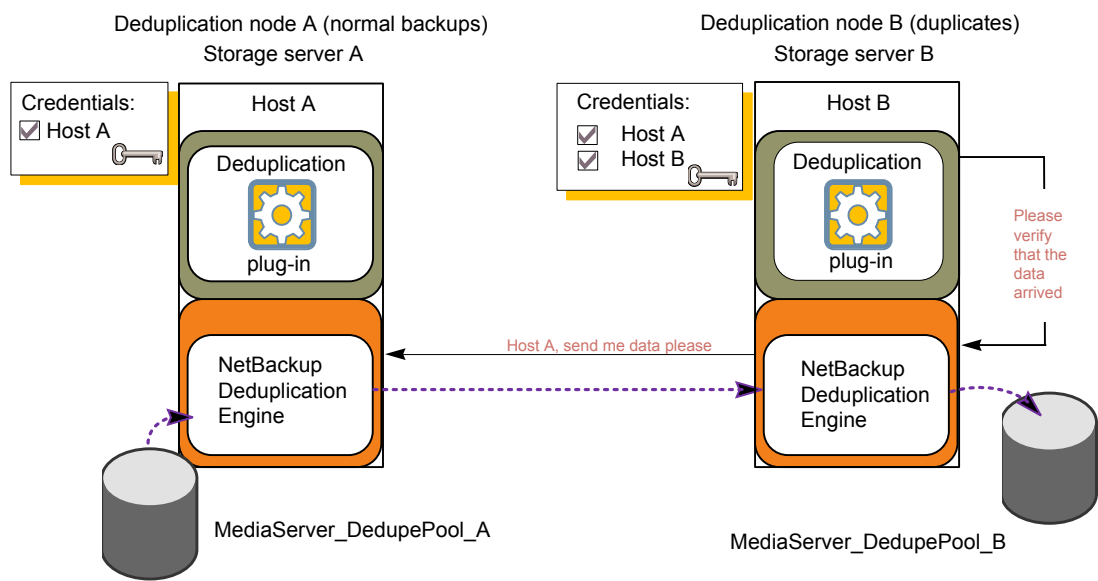


Figure 5-5 shows the storage unit settings for the duplication destination. They are similar to the push example except host B is selected. Host B is the common server, so it must be selected in the storage unit.

**Figure 5-5** Pull duplication storage unit settings

The screenshot shows the 'New Storage Unit' dialog box. The 'Storage unit name' field is filled with 'STU-Duplicates'. The 'Storage unit type' dropdown is set to 'Disk', and the 'On demand only' checkbox is checked. The 'Disk type' dropdown is set to 'PureDisk'. In the 'Properties' section, the 'Disk Pool' dropdown is set to 'MediaServer\_DedupePool\_B', with a 'Properties' button to its right. Under the 'Media Server' section, the radio button for 'Only use the following media servers' is selected. Below this, a list box titled 'Media Servers' contains two entries: 'Host\_A' (unchecked) and 'Host\_B' (checked).

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.

## Configuring MSDP optimized duplication within the same NetBackup domain

You can configure optimized duplication from a **Media Server Deduplication Pool** to other deduplication storage within the same NetBackup domain.

**Table 5-20** How to configure optimized duplication of deduplicated data

Step	Action	Description
Step 1	Review optimized duplication	See <a href="#">“About MSDP optimized duplication within the same domain”</a> on page 118.

**Table 5-20**      How to configure optimized duplication of deduplicated data  
(continued)

Step	Action	Description
Step 2	Configure the storage servers	<p>See <a href="#">"Configuring a storage server for a Media Server Deduplication Pool"</a> on page 67.</p> <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 <a href="#">"About the media servers for MSDP optimized duplication within the same domain"</a> on page 121.</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>
Step 3	Configure the deduplication pools	<p>If you did not configure the deduplication pools when you configured the storage servers, use the <b>Disk Pool Configuration Wizard</b> to configure them.</p> <p>See <a href="#">"Configuring a disk pool for deduplication"</a> on page 92.</p>
Step 4	Configure the storage unit for backups	<p>In the storage unit for your backups, do the following:</p> <ol style="list-style-type: none"> <li>For the <b>Disk type</b>, select <b>PureDisk</b>.</li> <li>For the <b>Disk pool</b>, select your <b>Media Server Deduplication Pool</b> or <b>PureDisk Deduplication Pool</b>.</li> </ol> <p><b>Note:</b> NetBackup supports <b>PureDisk Deduplication Pool</b> storage on NetBackup 5000 series appliances only.</p> <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 <a href="#">"Configuring a Media Server Deduplication Pool storage unit"</a> on page 102.</p>

**Table 5-20** How to configure optimized duplication of deduplicated data  
(continued)

Step	Action	Description
Step 5	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> <ol style="list-style-type: none"> <li>For the <b>Disk type</b>, select <b>PureDisk</b>.</li> <li>For the <b>Disk pool</b>, the destination can be a <b>Media Server Deduplication Pool</b> or a <b>PureDisk Deduplication Pool</b>.</li> </ol> <p><b>Note:</b> NetBackup supports <b>PureDisk Deduplication Pool</b> storage on NetBackup 5000 series appliances only.</p> <p>Also select <b>Only use the following media servers</b>. 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 <a href="#">“Configuring a Media Server Deduplication Pool storage unit”</a> on page 102.</p>
Step 6	Configure optimized duplication bandwidth	<p>Optionally, you can configure the bandwidth for replication.</p> <p>See <a href="#">“About configuring MSDP optimized duplication and replication bandwidth”</a> on page 151.</p>
Step 7	Configure optimized duplication behaviors	<p>Optionally, you can configure the optimized duplication behavior.</p> <p>See <a href="#">“Configuring NetBackup optimized duplication or replication behavior”</a> on page 130.</p> <p>See <a href="#">“About configuring MSDP optimized duplication and replication bandwidth”</a> on page 151.</p>



**Table 5-20** How to configure optimized duplication of deduplicated data  
(continued)

Step	Action	Description
Step 8	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"> <li>■ The first operation must be a <b>Backup</b>. For the <b>Storage</b> for the <b>Backup</b> operation, select the storage unit that is the target of your backups. That storage unit may use a <b>Media Server Deduplication Pool</b> or a <b>PureDisk Deduplication Pool</b>. These backups are the primary backup copies; they are the source images for the duplication operation.</li> <li>■ For the second, child <b>Operation</b>, select <b>Duplication</b>. Then, select the storage unit for the destination deduplication pool. That pool may be a <b>Media Server Deduplication Pool</b> or a <b>PureDisk Deduplication Pool</b>.</li> </ul> <p><b>Note:</b> NetBackup supports <b>PureDisk Deduplication Pool</b> storage on NetBackup 5000 series appliances only.</p> <p>See <a href="#">"About storage lifecycle policies"</a> on page 152.</p> <p>See <a href="#">"Creating a storage lifecycle policy"</a> on page 155.</p>
Step 9	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"> <li>■ If you use a storage lifecycle policy to manage the backup job and the duplication job: Select that storage lifecycle policy in the <b>Policy storage</b> field of the Policy <b>Attributes</b> tab.</li> <li>■ If you do not use a storage lifecycle policy to manage the backup job and the duplication job: Select the storage unit that contains your normal backups. These backups are the primary backup copies.</li> </ul> <p><b>Note:</b> NetBackup supports <b>PureDisk Deduplication Pool</b> storage on NetBackup 5000 series appliances only.</p> <p>See <a href="#">"About MSDP backup policy configuration"</a> on page 160.</p> <p>See <a href="#">"Creating a backup policy"</a> on page 161.</p>

**Table 5-20**      How to configure optimized duplication of deduplicated data  
(continued)

Step	Action	Description
Step 10	Configure NetBackup Vault for the duplication	<p>Configure Vault duplication only if you use NetBackup Vault to duplicate the images.</p> <p>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"> <li>Configure a Vault profile. <ul style="list-style-type: none"> <li>On the Vault <b>Profile</b> dialog box <b>Choose Backups</b> tab, choose the backup images in the source <b>Media Server Deduplication Pool</b> or a <b>PureDisk Deduplication Pool</b>.</li> <li>On the <b>Profile</b> dialog box <b>Duplication</b> tab, select the destination storage unit in the <b>Destination Storage unit</b> field.</li> </ul> </li> <li>Configure a Vault policy to schedule the duplication jobs. A Vault policy is a NetBackup policy that is configured to run Vault jobs.</li> </ul> <p><b>Note:</b> NetBackup supports <b>PureDisk Deduplication Pool</b> storage on NetBackup 5000 series appliances only.</p>
Step 11	Duplicate by using the <code>bpduplicate</code> command	<p>Use the NetBackup <code>bpduplicate</code> command only if you want to duplicate images manually.</p> <p>Duplicate from a <b>Media Server Deduplication Pool</b> or a <b>PureDisk Deduplication Pool</b> to another <b>Media Server Deduplication Pool</b> or another <b>PureDisk Deduplication Pool</b> in the same domain.</p> <p><b>Note:</b> NetBackup supports <b>PureDisk Deduplication Pool</b> storage on NetBackup 5000 series appliances only.</p> <p>See the <i>NetBackup Commands Reference Guide</i>:</p> <p><a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>

## Configuring NetBackup optimized duplication or replication behavior

You can configure some optimized duplication and replication behaviors for NetBackup. The behaviors depend on how NetBackup duplicates the images, as described in the following table.

**Table 5-21** Optimized duplication behavior

Behavior	Description
Duplication by using NetBackup Vault or the <code>bpduplicate</code> command	<p>If you use NetBackup Vault or the <code>bpduplicate</code> command for duplication, you can configure the following behaviors:</p> <ul style="list-style-type: none"> <li>Number of optimized duplication attempts. You can change the number of times NetBackup retries an optimized duplication job before it fails the jobs. See <a href="#">"To configure the number of duplication attempts"</a> on page 131.</li> <li>Optimized duplication failover. By default, if an optimized duplication job fails, NetBackup does not run the job again. You can configure NetBackup to use normal duplication if an optimized duplication job fails. See <a href="#">"To configure optimized duplication failover"</a> on page 132.</li> </ul>
Duplication or replication by using a storage lifecycle policy	<p>If a storage lifecycle policy optimized duplication or replication job fails, NetBackup waits 2 hours and retries the job again. NetBackup repeats the retry behavior until the job succeeds or the source backup image expire.</p> <p>You can change the number of hours for the wait period. See <a href="#">"To configure the storage lifecycle policy wait period"</a> on page 132.</p>

If you use a storage lifecycle policy for duplication, do not configure optimized duplication behavior for NetBackup Vault or the `bpduplicate` command, and vice versa. NetBackup behavior may not be predictable.

---

**Caution:** These settings affect all optimized duplication jobs; they are not limited to a specific NetBackup storage option.

---

#### To configure the number of duplication attempts

- ◆ On the master server, create a file named `OPT_DUP_BUSY_RETRY_LIMIT`. Add an integer to the file 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 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 the command line”](#) on page 132.

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

### To configure the storage lifecycle policy wait period

- 1 in the **NetBackup Administration Console**, expand **NetBackup Management** > **Host Properties** > **Master Servers**. Select the master server, then on the **Actions** menu click **Properties**.
- 2 Select **SLP Parameters**.
- 3 Change the **Extended image retry interval** to the new value.
- 4 Click **OK**.

## Setting NetBackup configuration options by using the command line

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

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

- `CLIENT_READ_TIMEOUT = 300`
- `OLD_VNETD_CALLBACK = YES`
- `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE`
- `SERVER = server1.NetBackup`

You can specify some options multiple times, such as the `SERVER` option.

Use the `bpgetconfig` command or `nbgetconfig` command to view your configuration information.

To set configuration options by using the command line

- 1
- In a command window or shell window on the host on which you want to set the property, invoke one of the following commands:

Windows NetBackup clients: `install_path\NetBackup\bin\nbsetconfig.exe`

Windows NetBackup servers:

`install_path\NetBackup\bin\admincmd\bpsetconfig.exe`

UNIX NetBackup clients: `/usr/opensv/netbackup/bin/nbsetconfig`

UNIX NetBackup servers: `/usr/opensv/netbackup/bin/admincmd/bpsetconfig`
- 2
- At the command 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:

Windows: `Ctrl + Z Enter`

UNIX: `Ctrl + D Enter`

# About MSDP replication to a different domain

NetBackup supports replication to storage in a different domain. NetBackup Auto Image Replication is the method used to replicate backup images. (Backup image replication is not the same as snapshot replication, which may occur in the same domain.) You can replicate from one source to one or more destinations.

[Table 5-22](#) describes the MSDP replication source and targets that NetBackup supports.

**Table 5-22** NetBackup media server deduplication replication targets

Source storage	Target storage
<b>Media Server Deduplication Pool</b> on a NetBackup media server	<b>A Media Server Deduplication Pool</b> on a NetBackup media server.

**Table 5-22** NetBackup media server deduplication replication targets  
(continued)

Source storage	Target storage
<b>Media Server Deduplication Pool</b> on a NetBackup media server	<p>A <b>PureDisk Deduplication Pool</b> on a NetBackup 5000 series appliance. A <b>PureDisk Deduplication Pool</b> is a PureDisk Storage Pool exposed to NetBackup by PureDisk.</p> <p>A <b>PureDisk Deduplication Pool</b> can be the destination of replication from one NetBackup domain only. If you try to replicate from MSDP storage in more than one NetBackup domain, replication does not function correctly.</p> <p><b>Note:</b> NetBackup supports <b>PureDisk Deduplication Pool</b> storage on NetBackup 5000 series appliances only.</p>
<b>Media Server Deduplication Pool</b> on a NetBackup 5200 series appliance	A <b>Media Server Deduplication Pool</b> on a NetBackup 5200 series appliance.

Auto Image Replication does not support replicating from a storage unit group. That is, the source copy cannot be in a storage unit group.

If a replication job fails, NetBackup retries the replication until it succeeds or the source images expire. You can change the retry interval behavior.

See [“Configuring NetBackup optimized duplication or replication behavior”](#) on page 130.

If a job fails after it replicates some of the images, NetBackup does not run a separate image cleanup job for the partially replicated images. The next time the replication runs, that job cleans up the image fragments before it begins to replicate the images.

You can use a separate network for the duplication traffic.

See [“About a separate network path for MSDP duplication and replication”](#) on page 116.

See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 135.

See [“About MSDP optimized duplication and replication”](#) on page 31.

# Configuring MSDP replication to a different NetBackup domain

Table 5-23 describes the tasks that are required to replicate backup images from one **Media Server Deduplication Pool** to another in a different NetBackup domain.

Optionally, you can use a separate network for the optimized duplication traffic.

See [“About a separate network path for MSDP duplication and replication”](#) on page 116.

Table 5-23 NetBackup MSDP replication configuration tasks

Step	Task	Procedure
Step 1	Learn about MSDP replication	See <a href="#">“About MSDP replication to a different domain”</a> on page 133. See <a href="#">“About NetBackup Auto Image Replication”</a> on page 136.
Step 2	Determine if you need to configure a trust relationship with the target NetBackup domain	A trust relationship is optional. See <a href="#">“About trusted master servers for Auto Image Replication”</a> on page 144.
Step 3	Add the remote storage server as a replication target	See <a href="#">“Configuring a target for MSDP replication to a remote domain”</a> on page 146. See <a href="#">“Viewing the replication topology for Auto Image Replication”</a> on page 142.

**Table 5-23** NetBackup MSDP replication configuration tasks (*continued*)

Step	Task	Procedure
Step 4	Configure a storage lifecycle policy	<p>The following are the options when you configure the SLP operations:</p> <ul style="list-style-type: none"> <li>■ If you configured a trust relationship with the target domains, you can specify one of the following options: <ul style="list-style-type: none"> <li>■ <b>All replication target storage servers (across different NetBackup domains)</b> NetBackup automatically creates an import SLP in the target domain when the replication job runs.</li> <li>■ <b>A specific Master Server.</b> If you choose this option, you then select <b>Target master server</b> and <b>Target import SLP</b>. You must create an import SLP in the target domain before you configure an SLP in the source domain.</li> </ul> </li> <li>■ If you did <i>not</i> configure a trust relationship with the target domains, <b>All replication target storage servers (across different NetBackup domains)</b> is selected by default. You cannot choose a specific target storage server. NetBackup automatically creates an import SLP in the target domain when the replication job runs.</li> </ul> <p>See <a href="#">“About storage lifecycle policies”</a> on page 152.</p> <p>See <a href="#">“About the storage lifecycle policies required for Auto Image Replication”</a> on page 153.</p> <p>See <a href="#">“Creating a storage lifecycle policy”</a> on page 155.</p>
Step 5	Configure replication bandwidth	<p>Optionally, you can configure the bandwidth for replication.</p> <p>See <a href="#">“About configuring MSDP optimized duplication and replication bandwidth”</a> on page 151.</p>

## About NetBackup Auto Image Replication

The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains. This process is referred to as Auto Image Replication.

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 138.



- 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.

Table 5-24 NetBackup Auto Image Replication storage types

Storage type	Link to more information
Media Server Deduplication Pool	See <a href="#">“Configuring MSDP replication to a different NetBackup domain”</a> on page 135.

Notes about Auto Image Replication

- Auto Image Replication is supported between master server domains that run any of the following versions of NetBackup:
  - NetBackup 7.1.0.4
  - NetBackup 7.5.0.3 and laterFor catalog backup images, NetBackup supports Auto Image Replication only between the same release levels of NetBackup.
- Although Auto Image Replication is a disaster recovery solution, you cannot restore to clients in the primary (or originating) domain from the target master domain.
- Auto Image Replication does not support synthetic backups.
- Auto Image Replication does not support replicating from a storage unit group. That is, the source copy cannot be in a storage unit group.
- Synchronize the clocks of the master servers in the source and the target domains so that the master server in the target domain can import the images as soon as they are ready. The master server in the target domain cannot import an image until the image creation time is reached. Time zone differences are not a factor because the images use Coordinated Universal Time (UTC).

Process Overview

[Table 5-25](#) is an overview of the process, generally describing the events in the originating and target domains.

NetBackup uses storage lifecycle policies in the source domain and the target domain to manage the Auto Image Replication operations.

See [“About the storage lifecycle policies required for Auto Image Replication”](#) on page 153.

**Table 5-25** Auto Image Replication process overview

Event	Domain in which event occurs	Event description
1	Originating master server (Domain 1)	Clients are backed up according to a backup policy that indicates a storage lifecycle policy as the <b>Policy storage</b> selection. The SLP must include at least one <b>Replication</b> operation to similar storage in the target domain.
2	Target master server (Domain 2)	The storage server in the target domain recognizes that a replication event has occurred. It notifies the NetBackup master server in the target domain.
3	Target master server (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 <b>Catalog</b> utility.)
4	Target master server (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.

## 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.

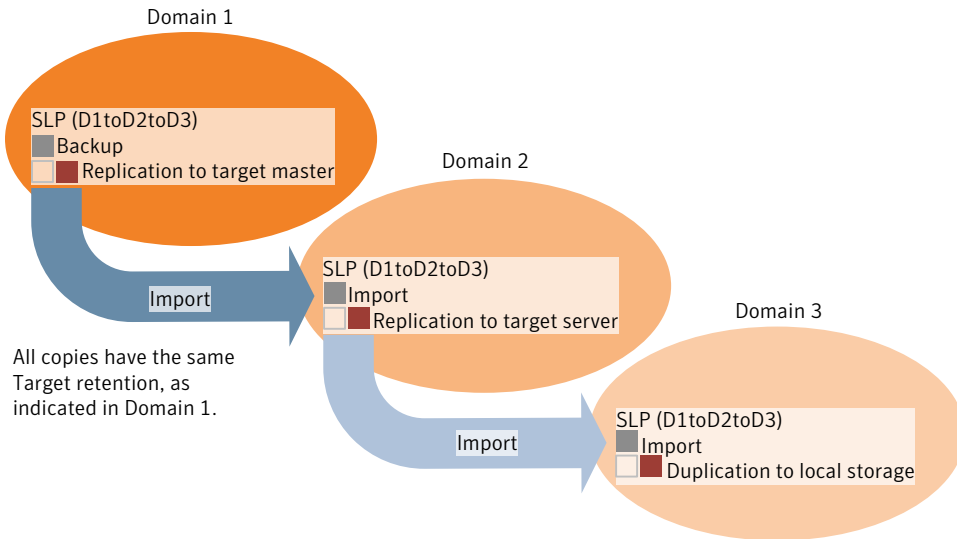
## Cascading Auto Image Replication model

Replications can be cascaded from the originating domain to multiple domains. 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 5-6](#) 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 5-6** 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 indicates that the master server in Domain 2 is the originating master server. However, after the image is imported successfully into Domain 3, NetBackup correctly indicates that the originating master server is in Domain 1.

---

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

The Import SLP must 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 5-26 shows the difference in the import operation setup.

**Table 5-26** 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 <b>Fixed</b> retention type	Same; no difference.
At least one operation must use the <b>Target retention</b> .	Here is the difference:  To meet the criteria, the import operation must use <b>Target retention</b> .

The target retention is embedded in the source image.

In the cascading model that is represented in [Figure 5-6](#), 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 5-7** Cascading replications to target master servers, with various target retentions



About the domain relationship for replication

For a **Media Server Deduplication Pool** target: the relationship between the originating domain and the target domain or domains is established in the originating domain. Specifically, by configuring the target storage server in the **Replication** tab of the **Change Storage Server** dialog box of the source storage server.

See “[Configuring a target for MSDP replication to a remote domain](#)” on page 146.

Before you configure the replication relationship, you can add the target master server as a trusted host.

See “[About trusted master servers for Auto Image Replication](#)” on page 144.

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

About the replication topology for Auto Image Replication

For Auto Image Replication, the disk volumes 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:

<b>Source</b>	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.
<b>Target</b>	A target volume in the remote domain is the replication partner of a source volume in the originating domain.
<b>None</b>	The volume does not have a replication attribute.

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

You configure the replication relationships in the source domain. To do so, you add target storage servers in the **Replication** tab of the **Change Storage Server** dialog box of the source storage server.

See “[Configuring a target for MSDP replication to a remote domain](#)” on page 146.

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

See [“Changing a Media Server Deduplication Pool properties”](#) on page 226.

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

- After you configure the replication targets.
- After you configure the storage server and before you configure disk pools.
- After changes to the volumes that comprise the storage.

See [“Viewing the replication topology for Auto Image Replication”](#) on page 142.

## Viewing the replication topology for Auto Image Replication

A volume that is a source of replication must have at least one replication partner that is the target of the replication. NetBackup lets you view the replication topology of the storage.

See [“About the replication topology for Auto Image Replication”](#) on page 141.

### 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:
  - Windows: `install_path\NetBackup\bin\admincmd\bpstsinfo -lsuinfo -storage_server host_name -stype server_type`
  - UNIX: `/usr/opensv/netbackup/bin/admincmd/bpstsinfo -lsuinfo -storage_server host_name -stype server_type`

The following are the options and arguments for the command:

<code>-storage_server host_name</code>	The name of the target storage server.
<code>-stype PureDisk</code>	Use <code>PureDisk</code> for a <b>Media Server Deduplication Pool</b> .

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

See [“Sample volume properties output for MSDP replication”](#) on page 142.

### 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_LSUF_DISK | STS_LSUF_ACTIVE | STS_LSUF_STORAGE_NOT_FREED |
    STS_LSUF_REP_ENABLED | STS_LSUF_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 trusted master servers for Auto Image Replication

NetBackup provides the ability to establish a trust relationship between replication domains. A trust relationship is optional. The following items describe how a trust relationship affects Auto Image Replication:

Trust relationship	You can select a subset of your trusted domains as a target for replication. NetBackup then replicates to the specified domains only rather than to all configured replication targets. This type of Auto Image Replication is known as Targeted AIR.
No trust relationship	NetBackup replicates to all defined target storage servers. You cannot select a specific host or hosts as a target.

You add a trusted master server in the source domain; you specify a remote master server as a trusted host. A trust relationship is reciprocal: If you add host B as a trusted master server for host A, host B then trusts hosts A.

You select the targets for replication when you configure a storage lifecycle policy. However, before you choose a specific storage server as a replication target, you must create an import SLP in the target domain. Then, you choose the specific target master server and SLP when you create a storage lifecycle policy in the source domain.

See [“About storage lifecycle policies”](#) on page 152.

See [“Creating a storage lifecycle policy”](#) on page 155.

---

**Note:** Before you can configure trust relationships for a clustered NetBackup master server, you must enable inter-node authentication on all of the nodes in the cluster. This requirement applies regardless of whether the clustered mastered server is the source of the replication operation or the target.

See [“Enabling NetBackup clustered master server inter-node authentication”](#) on page 145.

---



See [“About NetBackup Auto Image Replication”](#) on page 136.

## Enabling NetBackup clustered master server inter-node authentication

NetBackup requires inter-node authentication among the master servers in a cluster. For authentication, you must provision an authentication certificate on all of the nodes of the cluster. The certificates are used to establish SSL connections between the NetBackup hosts. The inter-node authentication allows the following NetBackup functionality:

NetBackup Administration Console	The NetBackup Administration Console in master server clusters requires the NetBackup authentication certificates for correct functionality.
Targeted AIR (Auto Image Replication)	<p>Auto Image Replication in which a master server is in a cluster requires inter-node authentication among the hosts in that cluster. The NetBackup authentication certificates provide the means to establish the proper trust relationships.</p> <p>Provision the certificates on the cluster hosts before you add the trusted master server. This requirement applies regardless of whether the clustered master server is the source of the replication operation or the target.</p> <p>See <a href="#">“About trusted master servers for Auto Image Replication”</a> on page 144.</p>

### To enable clustered master server inter-node authentication

- ◆ On the active node of the NetBackup master server cluster, run the following NetBackup command:

- **UNIX:** `/usr/opensv/netbackup/bin/admincmd/bpnbaz -setupat`
- **Windows:** `install_path\Veritas\NetBackup\bin\admincmd\bpnbaz -setupat`

NetBackup creates the certificates on every node in the master server cluster.

The following is example output:

```
# bpnbaz -setupat
You will have to restart NetBackup services on this machine after
the command completes successfully.
Do you want to continue(y/n)y
Gathering configuration information.
Please be patient as we wait for 10 sec for the security services
```

```
to start their operation.  
Generating identity for host 'bitl.remote.'  
Setting up security on target host: bitl.remote.  
nbatd is successfully configured on NetBackup Master Server.  
Operation completed successfully.
```

## Configuring a target for MSDP replication to a remote domain

Use the following procedure to configure a target for replication from a **Media Server Deduplication Pool** in an originating domain to a deduplication pool in another target domain. NetBackup supports several deduplication targets.

See [“About MSDP replication to a different domain”](#) on page 133.

Configuring the target storage server is only one step in the process of configuring MSDP replication.

See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 135.

---

**Note:** About clustered master servers: If you add a trusted master server for replication operations, you must enable inter-node authentication on all of the nodes in the cluster. Enable the authentication before you begin the following procedure. This requirement applies regardless of whether the clustered master server is the source of the replication operation or the target.

See [“About trusted master servers for Auto Image Replication”](#) on page 144.

See [“Enabling NetBackup clustered master server inter-node authentication”](#) on page 145.

---

---

**Caution:** Choose the target storage server or servers carefully. A target storage server must not also be a storage server for the source domain. Also, a disk volume must not be shared among multiple NetBackup domains.


---

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**.

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

The following is an example of the **Change Storage Server** dialog box **Replication** tab:



- 5 On the **Replication** tab, click **Add**. The **Add a Replication Target Across a Different NetBackup Domain** dialog box appears.

The following is an example of the dialog box.

**Add a Replication Target in a Different NetBackup Do...**

Replication target in a different NetBackup domain

Select trusted master server which manages the target storage device.  
Note: specifying a target master server is not mandatory. Add a target master server to setup a trust only if you want to replicate backups to specific master server domains.

Target master server: <Select target master server from the list>

Target storage server type: PureDisk

Target storage server name:

Target storage server details can be automatically fetched for trusted target master server to help in configuration.

Login credentials for the replication target storage server  
Required to establish an authenticated connection to the target storage server.

User name:

Password:

OK Cancel Help

- 6 In the **Add a Replication Target Across a Different NetBackup Domain** dialog box, complete one or more of the following procedures depending on your requirements:

Add a trusted master server

Add a trusted master server if you want to replicate backup images to a subset of available targets.

To add a trusted master server

**1** In the **Target master server** drop-down list, select **Add a new trusted master server**.

**2** Complete the fields in the **Add a New Trusted Master Server** dialog box. Click **OK** after you complete the fields.

See ["Target options for MSDP replication"](#) on page 150.

**3** Repeat the first two steps until you have finished adding trusted master servers.

**4** To add a replication target, continue with the next procedure.

See ["About trusted master servers for Auto Image Replication"](#) on page 144.

Select a trusted master server and replication target

Select a trusted master server if you want to replicate backup images to a subset of available targets.

To select a trusted master server and replication target

**1** In the **Target master server** drop-down list, select the master server of the domain to which you want to replicate data.

All trusted master servers are in the drop-down list.

**2** In the **Target storage server type** drop-down list, select the type of target storage server.

All available target types are in the drop-down list.

**3** In the **Target storage server name** drop-down list, select the storage server that hosts the target deduplication pool. All available storage servers in the target domain are in the drop-down list.

After you select the target storage server, NetBackup populates the **User name** field with the **User name** of the NetBackup Deduplication Engine of the target host.

**4** Enter the **Password** for the deduplication service on the target storage server.

**5** Click **OK**.

See ["Target options for MSDP replication"](#) on page 150.

Enter a replication target    Enter a replication target if you did not configure trusted master servers.

To enter a replication target

- 1    In the **Target storage server name** field, enter the name of the target storage server.
- 2    Enter the **User name** and **Password** for the NetBackup Deduplication Engine on the target storage server.
- 3    Click **OK**.

See [“Target options for MSDP replication”](#) on page 150.

- 7    After all replication targets are added, click **OK**.
- 8    For the deduplication pools 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, you must refresh the deduplication pools so that NetBackup reads the new volume properties.

See [“Changing a Media Server Deduplication Pool properties”](#) on page 226.

### Target options for MSDP replication

The following table describes the target options for replication to a NetBackup **Media Server Deduplication Pool**.

Table 5-27            MSDP target replication options

Option	Description
Target master server	All trusted master servers are in the drop-down list.  Select the master server for the target domain to which you want to replicate backups.  To add the master server of another domain as a trusted master, select <b>Add a new Trusted Master Server</b> . Configuring a trust relationship is required only if you want to choose a specific target for replication.

**Table 5-27** MSDP target replication options (*continued*)

Option	Description
<b>Target storage server type</b>	<p>The target storage server type must be <b>PureDisk</b>.</p> <p>If a trusted master server is configured, <b>PureDisk</b> is selected by default. If a trusted master server is <i>not</i> configured, NetBackup populates this field when it connects with the storage server you enter in the <b>Target storage server name</b> field.</p>
<b>Target storage server name</b>	<p>If a trusted master server is configured, select the target storage server. If a trusted master server is <i>not</i> configured, enter the name of the target storage server.</p> <p>The drop-down list shows all the storage servers that match the <b>Target storage server type</b>.</p>
<b>User name</b>	<p>When you configure a replication target, NetBackup populates the <b>User name</b> field with user account of the target storage server, as follows:</p> <ul style="list-style-type: none"> <li>For an MSDP target, the NetBackup Deduplication Engine user name.</li> <li>For a PureDisk storage target, the PureDisk Storage Pool Authority user name.</li> </ul> <p>NetBackup supports <b>PureDisk Deduplication Pool</b> storage on NetBackup 5000 series appliances only.</p>
<b>Password</b>	Enter the password for the NetBackup Deduplication Engine.

## About configuring MSDP 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:

**bandwidthlimit** The **bandwidthlimit** parameter in the `agent.cfg` file is the global bandwidth setting. You can use this parameter to limit the bandwidth that all replication jobs use. It applies to jobs in which a **Media Server Deduplication Pool** is the source. Therefore, configure it on the source storage server.

If **bandwidthlimit** is greater than zero, all of the jobs share the bandwidth. That is, the bandwidth for each job is the **bandwidthlimit** divided by the number of jobs.

If **bandwidthlimit=0**, total bandwidth is not limited. However, you can limit the bandwidth that each job uses. See the following **OPTDUP\_BANDWIDTH** description.

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

By default, **bandwidthlimit=0**.

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

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

**OPTDUP\_BANDWIDTH** The **OPTDUP\_BANDWIDTH** parameter in the `pd.conf` file specifies the per job bandwidth.

**OPTDUP\_BANDWIDTH** applies only if the **bandwidthlimit** parameter in the `agent.cfg` file is zero.

If **OPTDUP\_BANDWIDTH** and **bandwidthlimit** are both 0, bandwidth per replication job is not limited.

By default, **OPTDUP\_BANDWIDTH = 0**.

See [“About the MSDP pd.conf configuration file”](#) on page 169.

See [“Editing the MSDP pd.conf file”](#) on page 170.

See [“MSDP pd.conf file parameters”](#) on page 171.

See [“Configuring MSDP optimized duplication within the same NetBackup domain”](#) on page 126.

See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 135.

## About storage lifecycle policies

A storage lifecycle policy (SLP) is a storage plan for a set of backups. An SLP is configured within the **Storage Lifecycle Policies** utility.



An SLP contains instructions in the form of storage operations, to be applied to the data that is backed up by a backup policy. Operations are added to the SLP that determine how the data is stored, copied, replicated, and retained. NetBackup retries the copies as necessary to ensure that all copies are created.

SLPs offer the opportunity for users to assign a classification to the data at the policy level. A data classification represents a set of backup requirements, which makes it easier to configure backups for data with different requirements. For example, email data and financial data.

SLPs can be set up to provide staged backup behavior. They simplify data management by applying a prescribed behavior to all the backup images that are included in the SLP. This process allows the NetBackup administrator to leverage the advantages of disk-based backups in the near term. It also preserves the advantages of tape-based backups for long-term storage.

The **SLP Parameters** properties in the **NetBackup Administration Console** allow administrators to customize how SLPs are maintained and how SLP jobs run.

Best-practice information about SLPs appears in the following document:

<http://www.symantec.com/docs/TECH208536>

For more information, see the [NetBackup Administrator's Guide, Volume I](#).

## About the storage lifecycle policies required for Auto Image Replication

To replicate images from one NetBackup domain to another NetBackup domain requires two storage lifecycle policies. The following table describes the policies and their requirements:

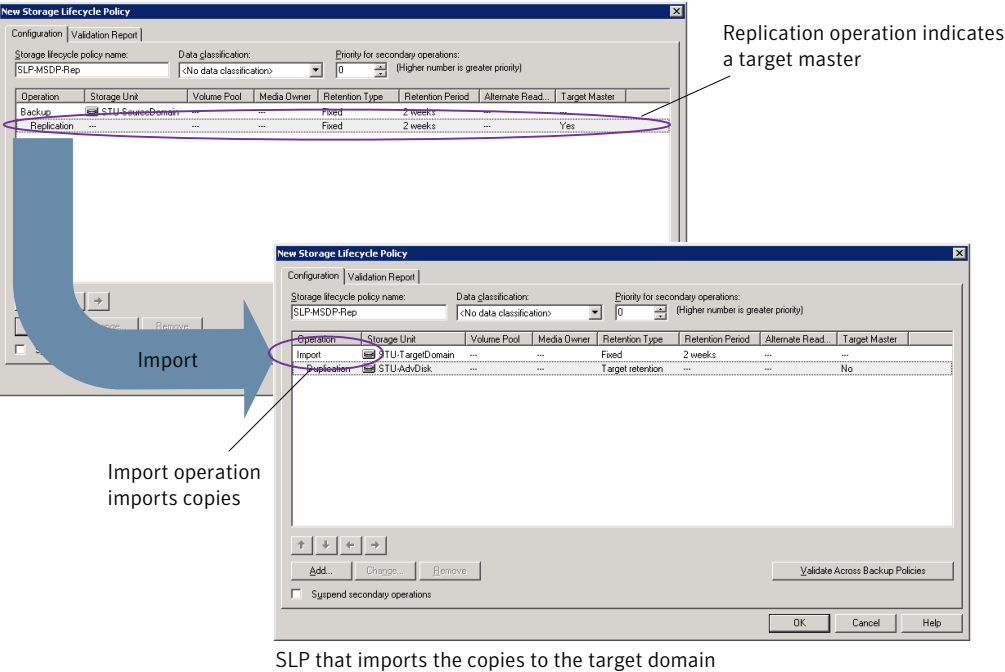
**Table 5-28** SLP requirements for Auto Image Replication

Domain	Storage lifecycle policy requirements
Domain 1 (Source domain)	<p>The Auto Image Replication SLP in the source domain must meet the following criteria:</p> <ul style="list-style-type: none"> <li>■ The first operation must be a <b>Backup</b> operation to a <b>Media Server Deduplication Pool</b>. Indicate the exact storage unit from the drop-down list. Do not select <b>Any Available</b>. <b>Note:</b> The target domain must contain the same type of storage to import the image.</li> <li>■ At least one operation must be a <b>Replication</b> operation to a <b>Media Server Deduplication Pool</b> in another NetBackup domain. You can configure multiple <b>Replication</b> operations in an Auto Image Replication SLP. The <b>Replication</b> operation settings determine whether the backup is replicated to all replication targets in all master server domains or only to specific replication targets. See <a href="#">“About trusted master servers for Auto Image Replication”</a> on page 144.</li> <li>■ The SLP must be of the same data classification as the <b>Import</b> SLP in Domain 2. <b>Note:</b> If the source master server is at a NetBackup version earlier than 7.6 and the master server in the target domain is at NetBackup 7.6: Do not use the data classification of <b>Any</b>. Use a different data classification in the source domain or the <b>Import</b> job fails.</li> </ul>
Domain 2 (Target domain)	<p>If replicating to all targets in all domains, in each domain NetBackup automatically creates an <b>Import</b> SLP that meets all the necessary criteria.</p> <p><b>Note:</b> If replicating to specific targets, you must create the <b>Import</b> SLP before creating the Auto Image Replication SLP in the originating domain.</p> <p>The <b>Import</b> SLP must meet the following criteria:</p> <ul style="list-style-type: none"> <li>■ The first operation in the SLP must be an <b>Import</b> operation. NetBackup must support the <b>Destination storage</b> as a target for replication from the source storage. Indicate the exact storage unit from the drop-down list. Do not select <b>Any Available</b>.</li> <li>■ The SLP must contain at least one operation that has the <b>Target retention</b> specified.</li> <li>■ The SLP must be of the same data classification as the SLP in Domain 1. Matching the data classification keeps a consistent meaning to the classification and facilitates global reporting by data classification.</li> </ul> <p>See the following topic for more information about <b>Replication</b> operation configuration:</p>

[Figure 5-8](#) shows how the SLP in the target domain is set up to replicate the images from the originating master server domain.

Figure 5-8 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

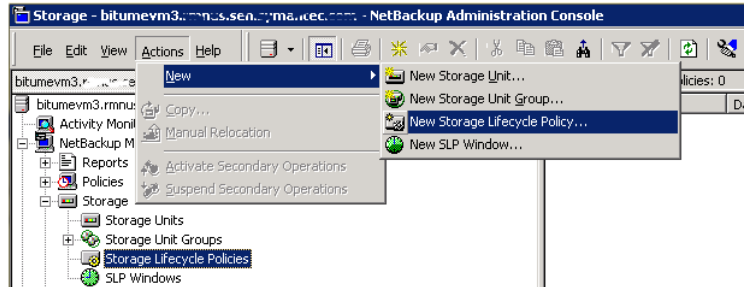
**Note:** Restart `nbstserv` after you make changes to the underlying storage for any operation in an SLP.

# Creating a storage lifecycle policy

A storage lifecycle policy (SLP) is a storage plan for a set of backups. The operations in an SLP are the backup instructions for the data. Use the following procedure to create an SLP that contains multiple storage operations.

To add a storage operation to a storage 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 In the **New Storage Lifecycle Policy** dialog box, enter a **Storage lifecycle policy name**.


See [“NetBackup naming conventions”](#) on page 19.

- 4 Add one or more operations to the SLP. The operations are the instructions for the SLP to follow and apply to the data that is specified in the backup policy.


If this is the first operation added to the SLP, click **Add**.

If this is not the first operation in the SLP, add an operation that is either hierarchical or non-hierarchical:

To create a hierarchical operation, select an operation to become the source of the next operation. Click **Add**. The new operation is a child of the selected operation. The child is indented under the parent operation.



To create a non-hierarchical operation, do not select an operation. A non-hierarchical operation means that the operation does not have a parent and child relationship with another operation. The new operation is not indented.



- 5 In the **Properties** tab of the **New Storage Operation** dialog box, select an **Operation** type. If you're creating a child operation, the SLP displays only those operations that are valid based on the parent operation that you've selected.

The name of the operation reflects its purpose in the SLP:

- **Backup**
- **Duplication**
- **Import**
- **Replication**

See [“About NetBackup Auto Image Replication”](#) on page 136.

- 6 Configure the properties for the operation.

- 7 The **Window** tab displays for the following operation types: **Backup From Snapshot**, **Duplication**, **Import**, **Index From Snapshot**, and **Replication**. If you'd like to control when the secondary operation runs, create a window for the operation.
- 8 Click the **Advanced** button in the **Properties** tab to display options about how the window should behave if the window closes and a secondary operation is not yet complete.
- 9 Click **OK** to create the operation.
- 10 Add additional operations to the SLP as needed. (See step 4.)
- 11 Change the hierarchy of the operations in the SLP if necessary.
- 12 Click **OK** to create the SLP. NetBackup validates the SLP when it is first created and whenever it is changed.
- 13 Configure a backup policy and select a storage lifecycle policy as the **Policy storage**.

See “[Creating a backup policy](#)” on page 161.

## Storage Lifecycle Policy dialog box settings

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

Figure 5-9 Storage Lifecycle Policy tab

The screenshot shows the 'Change Storage Lifecycle Policy' dialog box with the 'Storage Lifecycle Policy' tab selected. The 'Validation Report' sub-tab is also visible. The 'Storage lifecycle policy name' is 'SLP\_1\_snapshot'. The 'Data classification' is '<No data classification>'. The 'Priority for secondary operations' is '0'. Below these fields is a table with columns: Operation, Window, Storage, Volume Pool, Media Owner, Retention Type, Retention Peri..., Alternate Res..., and Target Master. The table contains two rows: 'Snapshot' and 'Backup From Sn...'. Below the table are buttons for 'Add...', 'Change...', and 'Remove'. At the bottom, there are radio buttons for 'Active' and 'Postponed', a checkbox for 'Urgent', and a 'Validate Across Backup Policies' button. The 'OK', 'Cancel', and 'Help' buttons are at the very bottom.

Operation	Window	Storage	Volume Pool	Media Owner	Retention Type	Retention Peri...	Alternate Res...	Target Master
Snapshot	---	No Storage Unit	---	---	Maximum snap...	---	---	---
Backup From Sn...	Default_24c7...	stu_adv_ca...	---	---	Fixed	2 weeks	---	---

Table 5-29 Storage Lifecycle Policy tab

Setting	Description
Storage lifecycle policy name	The <b>Storage lifecycle policy name</b> describes the SLP. The name cannot be modified after the SLP is created.
Data classification	<p>The <b>Data classification</b> defines the level or classification of data that the SLP is allowed to process. The drop-down menu contains all of the defined classifications as well as the <b>Any</b> classification, which is unique to SLPs.</p> <p>The <b>Any</b> selection indicates to the SLP that it should preserve all images that are submitted, regardless of their data classification. It is available for SLP configuration only and is not available to configure a backup policy.</p> <p>In an Auto Image Replication configuration where the master server domains run different versions of NetBackup, see the following topic for special considerations:</p> <p>See <a href="#">“About the storage lifecycle policies required for Auto Image Replication”</a> on page 153.</p> <p>The <b>Data classification</b> is an optional setting.</p> <p>One data classification can be assigned to each SLP and applies to all operations in the SLP.</p> <p>If a data classification is selected (other than <b>Any</b>), 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 <b>Data classification</b> 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>
Priority for secondary operations	<p>The <b>Priority for secondary operations</b> option is the priority that jobs from secondary operations have in relationship to all other jobs. The priority applies to the jobs that result from all operations except for <b>Backup</b> and <b>Snapshot</b> operations. Range: 0 (default) to 99999 (highest priority).</p> <p>For example, you may want to set the <b>Priority for secondary operations</b> for a policy with a gold data classification 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 <b>Attributes</b> tab.</p>

Table 5-29 Storage Lifecycle Policy tab (*continued*)

Setting	Description
<b>Operations</b>	<p>Use the <b>Add</b>, <b>Change</b>, and <b>Remove</b> buttons to create a list of operations in the SLP. An SLP must contain one or more operations. Multiple operations imply that multiple copies are created.</p> <p>The list also contains the columns that display information about each operation. Not all columns display by default.</p> <p>For column descriptions, see the following topic:</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>
<b>Active</b> and <b>Postponed</b>	<p>The <b>Active</b> and <b>Postponed</b> options appear under <b>State of Secondary Operation Processing</b> and refer to the processing of all duplication operations in the SLP.</p> <p><b>Note:</b> The <b>Active</b> and <b>Postponed</b> options apply to duplication operations that create tar-formatted images. For example, those created with <code>bpduplicate</code>. The <b>Active</b> and <b>Postponed</b> options do not affect the images that are duplicated as a result of OpenStorage optimized duplication, NDMP, or if one or more destination storage units are specified as part of a storage unit group.</p> <p>These options do not apply if either the source media server or the destination media server is not at NetBackup 7.6 or later.</p> <ul style="list-style-type: none"><li>■ Enable <b>Active</b> to let secondary operations continue as soon as possible. When changed from <b>Postponed</b> to <b>Active</b>, NetBackup continues to process the images, picking up where it left off when secondary operations were made inactive.</li><li>■ Enable <b>Postponed</b> to postpone the secondary operations for the entire SLP. <b>Postponed</b> does not postpone the creation of duplication jobs, it postpones the creation of images instead. The duplication jobs continue to be created, but they are not run until secondary operations are active again.</li></ul> <p>All secondary operations in the SLP are inactive indefinitely unless the administrator selects <b>Active</b> or until the <b>Until</b> option is selected and an activation date is indicated.</p>
<b>Validate Across Backup Policies</b> button	<p>Click 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 <b>Validation Report</b> tab.</p> <p>This button performs the same validation as the <code>-conflict</code> option performs when used with the <code>nbsl</code> command.</p>

## About MSDP 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 backup policy

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.

---

**Note:** Do not use the Policy Configuration Wizard to configure policies for Replication Director.

---

### Using the Policy Configuration Wizard to create a backup policy

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

To create a backup 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, 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 backup policy without using the Policy Configuration Wizard

Use the following procedure to create a backup policy in the **NetBackup Administration Console** 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 On the **Actions** menu, click **New > Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box.  
See [“NetBackup naming conventions”](#) on page 19.
- 4 If necessary, clear the **Use Policy Configuration Wizard** check box.
- 5 Click **OK**.
- 6 Configure the attributes, the schedules, the clients, and the backup selections for the new policy.

## 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 for backups and restores. A resilient connection allows backup and restore traffic between a client and a NetBackup media server to function effectively in high-latency, low-bandwidth networks such as WANs. 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.

NetBackup protects only the network socket connections that the NetBackup Remote Network Transport Service (`nbbrntd`) creates. Examples of the connections that are not supported are:

- Clients that back up their own data (deduplication clients and SAN clients)

- Granular Recovery Technology (GRT) for Exchange Server or SharePoint Server
- NetBackup `nbfssd` process.

NetBackup protects connections only after they are established. If NetBackup cannot create a connection because of network problems, there is nothing to protect.

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 5-10** Master server Resilient Network host properties

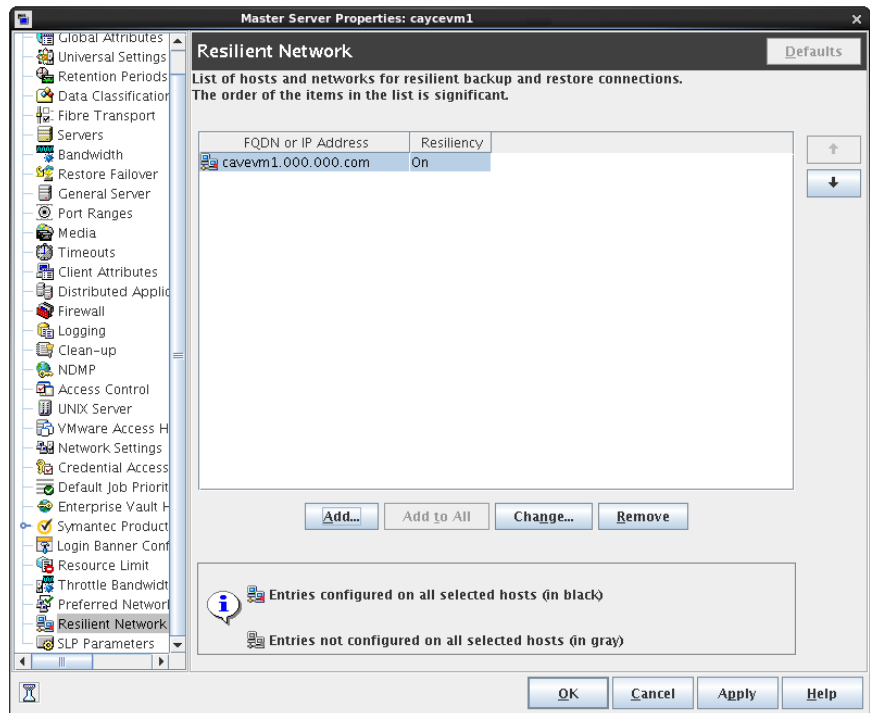


Table 5-30 describes the **Resilient Network** properties.

**Table 5-30** Resilient Network dialog box properties

Property	Description
<b>Host Name or IP Address</b>	<p>The <b>Host Name or IP Address</b> 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>
<b>Resiliency</b>	<b>Resiliency</b> is either <b>ON</b> or <b>OFF</b> .

---

**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.

---

**Note:** If multiple backup streams run concurrently, the Remote Network Transport Service writes a large amount of information to the log files. In such a scenario, Symantec recommends that you set the logging level for the Remote Network Transport Service to 2 or less. Instructions to configure unified logs are in a different guide.

See the [NetBackup Troubleshooting Guide](#).

---

## 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 162.

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

- Windows:  
`install_path\Veritas\NetBackup\bin\admincmd\resilient_clients`
- UNIX: `/usr/opensv/netbackup/bin/admincmd/resilient_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 master server on which to specify properties.
- 3 On the **Actions** menu, click **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 for clients:

**Add**

To add resilient settings

- 1 Click **Add**.

The **Add Resilient Network Settings** dialog box appears

- 2 Enter a client host name, an IP address, or an address range.

If you specify the client host by name, Symantec recommends that you use the fully qualified domain name.

For address ranges, use Classless Inter-Domain Routing (CIDR) notation (for example, **192.168.100.0/24** or **fd00::/8**).

- 3 Ensure that the **Resiliency On** option is selected.
- 4 Click **Add**.
- 5 Repeat until you have finished entering clients or address ranges.
- 6 When you finish adding network settings, click **Close**.

**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**

To change resilient settings

- 1 Select the client host name, the IP address, or the address range.
- 2 Click **Change**.

The **Change Resilient Network Settings** dialog box appears

- 3 Select the desired **Resiliency** setting.
- 4 Click **OK**.

**Remove**

Remove the select host or address rang

- 1 Select the client host name, the IP address, or the address range.
- 2 Click **Remove**.

The client is removed immediately; a confirmation dialog box does not appear.



Move an item in the list of items

- 1 Select the client host name, the IP address, or the address range.
- 2 Click the appropriate button to move up the item or move down the item.

The order of the items in the list is significant.

See [“Resilient Network properties”](#) on page 162.

- 6 After specifying resilient connections, click **OK**.

The settings are propagated to the affected hosts through normal NetBackup inter-host communication, which can take up to 15 minutes.

- 7 If you want to begin a backup immediately, restart the NetBackup services on the master server.

## Adding an MSDP load balancing server

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


See [“About MSDP storage servers”](#) on page 21.

**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 MSDP `pd.conf` configuration file

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 the 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`

See [“MSDP pd.conf file parameters”](#) on page 171.

See [“Editing the MSDP pd.conf file”](#) on page 170.

## Editing the MSDP pd.conf 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.

---

**Note:** Symantec recommends that you make a backup copy of the file before you edit it.

---

See [“About the MSDP pd.conf configuration file”](#) on page 169.

See [“MSDP pd.conf file parameters”](#) on page 171.

### 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.

## MSDP pd.conf file parameters

[Table 5-31](#) describes the deduplication parameters that you can configure for a NetBackup **Media Server Deduplication Pool** environment.

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

The parameters in the file in your release may differ from those that are described in this topic.

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

**Table 5-31** `pd.conf` file parameters

Parameter	Description
BACKUPRESTORERANGE	<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"><li>■ 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></li><li>■ 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></li></ul> <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>

Table 5-31 pd.conf file parameters (continued)

Parameter	Description
BANDWIDTH_LIMIT	<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>
COMPRESSION	<p>Specifies whether to compress the data during backups.</p> <p>By default, the data is compressed.</p> <p>Default value: <code>COMPRESSION = 1</code></p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See <a href="#">“About MSDP compression”</a> on page 108.</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.</p> <p>Consider the following information before you change this setting:</p> <ul style="list-style-type: none"><li>■ If disabled (set to 0), a request for the latest storage capacity information occurs whenever NetBackup requests it.</li><li>■ 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.</li><li>■ 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.</li><li>■ On high load systems, the load may delay the capacity information reporting. If so, NetBackup may mark the storage unit as down.</li></ul> <p>Default value: <code>CR_STATS_TIMER = 0</code></p> <p>Possible values: 0 or greater, in seconds</p>

Table 5-31 pd.conf file parameters (continued)

Parameter	Description
DEBUGLOG	<p>Specifies the file to which NetBackup writes the deduplication plug-in 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><b>Note:</b> This parameter does not apply for NDMP backups from a NetApp appliance.</p> <p>Default value:</p> <ul style="list-style-type: none"> <li>■ UNIX: DEBUGLOG = /var/log/puredisk/pdplugin.log</li> <li>■ Windows: DEBUGLOG = C:\pdplugin.log</li> </ul> <p>Possible values: Any path</p>
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. Note: this parameter does not apply to the NDMP backups that use the NetApp stream handler.</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 during backups. 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 <a href="#">“About MSDP encryption”</a> on page 110.</p>
FIBRECHANNEL	<p>Enable Fibre Channel for backup and restore traffic to and from a NetBackup series appliance.</p> <p>See <a href="#">“About Fibre Channel to a NetBackup 5020 appliance”</a> on page 308.</p> <p>Default value: FIBRECHANNEL = 0</p> <p>Possible values: 0 (off) or 1 (on)</p>

Table 5-31 pd.conf file parameters (*continued*)

Parameter	Description
FILE_KEEP_ALIVE_INTERVAL	<p>The interval in seconds at which to perform keepalives on idle sockets.</p> <p>The following items describe the behavior based on how you configure this parameter:</p> <ul style="list-style-type: none"><li>■ Commented out (default) and Resilient Network connections are enabled: If the value is less than 75 seconds, the keep alive interval is 60 seconds. If the value is greater than 1800 seconds (30 minutes), the keep alive interval is 1440 seconds (80% of 30 minutes). If the value is between 75 and 1800 sections, the keep-alive interval is 80% of the parameter value. See <a href="#">“Resilient Network properties”</a> on page 162.</li><li>■ Commented out (the default) and Resilient Network connections are <i>not</i> enabled: The keep-alive interval is 1440 seconds (80% of 30 minutes).</li><li>■ 0 or less: Disabled; no keepalives are sent.</li><li>■ Greater than 0: The keep-alive interval is the specified value in seconds except as follows: If less than 60 seconds or greater than 7200 seconds (two hours), the keep-alive interval is 1440 seconds (80% of 30 minutes).</li></ul> <p>Default value : <code>FILE_KEEP_ALIVE_INTERVAL = 1440</code></p> <p>Possible values: 0 (disabled) or 60 to 7200 seconds</p> <p>To determine the keep alive interval that NetBackup uses, examine the deduplication plug-in log file for a message similar to the following:</p> <p>Using keepalive interval of xxxx seconds</p> <p>For more information about the deduplication plug-in log file, see <code>DEBUGLOG</code> and <code>LOGLEVEL</code> in this table.</p>

Table 5-31 pd.conf file parameters (continued)

Parameter	Description
FP_CACHE_CLIENT_POLICY	<p><b>Note:</b> 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><b><i>clienthostmachine,backuppolicy,date</i></b></p> <p>The <b><i>date</i></b> 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 <a href="#">“Configuring MSDP fingerprint cache seeding on the client”</a> on page 63.</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><b>Note:</b> Change this value only when directed to do so by a Symantec representative.</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 = 1</p> <p>Possible values: 0 (off) or 1 (on)</p>

Table 5-31 pd.conf file parameters (continued)

Parameter	Description
FP_CACHE_MAX_COUNT	<p>Specifies the maximum number of images to load in the fingerprint cache.</p> <p>Default value: <code>FP_CACHE_MAX_COUNT = 1024</code></p> <p>Possible values: 0 to 4096</p> <p><b>Note:</b> Change this value only when directed to do so by a Symantec representative.</p>
FP_CACHE_MAX_MBSIZE	<p>Specifies the amount of memory in MBs to use for the fingerprint cache.</p> <p>Default value: <code>FP_CACHE_MAX_MBSIZE = 20</code></p> <p>Possible values: 0 to the computer limit</p> <p><b>Note:</b> Change this value only when directed to do so by a Symantec representative.</p>
FP_CACHE_PERIOD_REBASING_THRESHOLD	<p>Specifies the threshold (MB) for periodic rebasing during backups. A container is considered for rebasing if both of the following are true:</p> <ul style="list-style-type: none"><li>■ The container has not been rebased within the last three months.</li><li>■ For that backup, the data segments in the container consume less space than the <code>FP_CACHE_PERIOD_REBASING_THRESHOLD</code> value.</li></ul> <p>Default value: <code>FP_CACHE_PERIOD_REBASING_THRESHOLD = 16</code></p> <p>Possible values: 0 (disabled) to 256</p> <p>See <a href="#">“About MSDP storage rebasing”</a> on page 241.</p>
FP_CACHE_REBASING_THRESHOLD	<p>Specifies the threshold (MB) for normal rebasing during backups. A container is considered for rebasing if both of the following are true:</p> <ul style="list-style-type: none"><li>■ The container has been rebased within the last three months.</li><li>■ For that backup, the data segments in the container consume less space than the <code>FP_CACHE_REBASING_THRESHOLD</code> value.</li></ul> <p>Default value: <code>FP_CACHE_REBASING_THRESHOLD = 4</code></p> <p>Possible values: 0 (disabled) to 200</p> <p>If you change this value, consider the new value carefully. If you set it too large, all containers become eligible for rebasing. Deduplication rates are lower for the backup jobs that perform rebasing.</p> <p>See <a href="#">“About MSDP storage rebasing”</a> on page 241.</p>



Table 5-31 pd.conf file parameters (continued)

Parameter	Description
LOCAL_SETTINGS	<p>Specifies whether to use the <code>pd.conf</code> settings of the local host or to allow the server to override the local settings. The following is the order of precedence for local settings:</p> <ul style="list-style-type: none"><li>■ Local host</li><li>■ Load balancing server</li><li>■ Storage server</li></ul> <p>To use the local settings, set this value to 1.</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><b>Note:</b> 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><b>Note:</b> 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 = 100</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><b>Note:</b> Change this value only when directed to do so by a Symantec representative.</p>

Table 5-31 pd.conf file parameters (continued)

Parameter	Description
MTSTRM_BACKUP_CLIENTS	<p>If set, limits the use of the Multi-Threaded Agent to the backups of the specified clients. The clients that are not specified use single-threading.</p> <p>This setting does not guarantee that the specified clients use the Multi-Threaded Agent. The <code>MaxConcurrentSessions</code> parameter in the <code>mtstrm.conf</code> file controls the number of backups the Multi-Threaded Agent processes concurrently. If you specify more clients than the <code>MaxConcurrentSessions</code> value, some of the clients may use single-threaded processing.</p> <p>See <a href="#">“MSDP mtstrm.conf file parameters”</a> on page 55.</p> <p>The format is a comma-separated list of the clients, case insensitive (for example, <code>MTSTRM_BACKUP_CLIENTS = client1,client2,client3</code>).</p> <p>Default value: <code>MTSTRM_BACKUP_CLIENTS = (no default value)</code></p> <p>Possible values: comma separated client names</p> <p>See <a href="#">“About the MSDP Deduplication Multi-Threaded Agent”</a> on page 52.</p>

Table 5-31 pd.conf file parameters (continued)

Parameter	Description
MTSTRM_BACKUP_ENABLED	<p>Use the Multi-Threaded Agent in the backup stream between the deduplication plug-in and the NetBackup Deduplication Engine.</p> <p>Default value: <code>MTSTRM_BACKUP_ENABLED = (no default value)</code></p> <p>Possible values: 1 (On) or 0 (Off)</p> <p>The value for this parameter is configured during installation or upgrade. If the hardware concurrency value of the host is greater than a hardware concurrency threshold value, NetBackup sets <code>MTSTRM_BACKUP_ENABLED</code> to 1. (For the purposes of this parameter, the <i>hardware concurrency</i> is the number of CPUs or cores or hyperthreading units.)</p> <p>The following items describe the values that are used for the determination algorithm:</p> <ul style="list-style-type: none"><li>■ The hardware concurrency value is one of the following:<ul style="list-style-type: none"><li>■ For media servers, half of the host's hardware concurrency is used for the hardware concurrency value in the algorithm.</li><li>■ For clients, all of the host's hardware concurrency is used for the hardware concurrency value in the algorithm.</li></ul></li><li>■ The hardware concurrency threshold value to enable multithreading is one of the following:<ul style="list-style-type: none"><li>■ Windows and Linux: The threshold value is 2.</li><li>■ Solaris: The threshold value is 4.</li></ul></li></ul> <p>The following examples may be helpful:</p> <ul style="list-style-type: none"><li>■ A Linux media server that has 8 CPU cores with two hyperthreading units per core has a hardware concurrency of 16. Therefore, the hardware concurrency value for the algorithm is 8 (for media servers, half of the system's hardware concurrency). Eight is greater than two (the threshold value of Windows and Linux), so multithreading is enabled (<code>MTSTRM_BACKUP_ENABLED = 1</code>).</li><li>■ A Solaris client that has 2 CPU cores without hyperthreading has a hardware concurrency of 2. The hardware concurrency value for the algorithm is 2 (for clients, all of the system's hardware concurrency). Two is not greater than four (the threshold value of Solaris), so multithreading is not enabled (<code>MTSTRM_BACKUP_ENABLED = 0</code>).</li></ul> <p>See <a href="#">"About the MSDP Deduplication Multi-Threaded Agent"</a> on page 52.</p>

Table 5-31 pd.conf file parameters (continued)

Parameter	Description
MTSTRM_BACKUP_POLICIES	<p>If set, limits the use of the Multi-Threaded Agent to the backups of the specified policies. The clients in the policies that are not specified use single-threading, unless the client is specified in the <code>MTSTRM_BACKUP_CLIENTS</code> parameter.</p> <p>This setting does not guarantee that all of the clients in the specified policies use the Multi-Threaded Agent. The <code>MaxConcurrentSessions</code> parameter in the <code>mtstrm.conf</code> file controls the number of backups the Multi-Threaded Agent processes concurrently. If the policies include more clients than the <code>MaxConcurrentSessions</code> value, some of the clients may use single-threaded processing.</p> <p>See <a href="#">“MSDP mtstrm.conf file parameters”</a> on page 55.</p> <p>The format is a comma-separated list of the policies, case sensitive (for example, <code>MTSTRM_BACKUP_POLICIES = policy1,policy2,policy3</code>).</p> <p>Default value: <code>MTSTRM_BACKUP_POLICIES = (no default value)</code></p> <p>Possible values: comma separated backup policy names</p> <p>See <a href="#">“About the MSDP Deduplication Multi-Threaded Agent”</a> on page 52.</p>
MTSTRM_IPC_TIMEOUT	<p>The number of seconds to wait for responses from the Multi-Threaded Agent before the deduplication plug-in times out with an error.</p> <p>Default value: <code>MTSTRM_IPC_TIMEOUT = 1200</code></p> <p>Possible values: 1-86400, inclusive</p> <p>See <a href="#">“About the MSDP Deduplication Multi-Threaded Agent”</a> on page 52.</p>
OPTDUP_BANDWIDTH	<p>Determines the bandwidth that is allowed for each optimized duplication and Auto Image Replication stream on a deduplication server. <code>OPTDUP_BANDWIDTH</code> does not apply to clients. The value is specified in KBytes/second.</p> <p>Default value: <code>OPTDUP_BANDWIDTH= 0</code></p> <p>Possible values: 0 (no limit) to the practical system limit, in KBs/sec</p> <p>A global bandwidth parameter effects whether or not <code>OPTDUP_BANDWIDTH</code> applies.</p> <p>See <a href="#">“About configuring MSDP optimized duplication and replication bandwidth”</a> on page 151.</p>

**Table 5-31** pd.conf file parameters (*continued*)

Parameter	Description
OPTDUP_COMPRESSION	<p>Specifies whether to compress the data during optimized duplication and Auto Image Replication. By default, files are compressed. To disable compression, change the value to 0. This parameter does not apply to clients.</p> <p>Default value: OPTDUP_COMPRESSION = 1</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See <a href="#">“About MSDP compression”</a> on page 108.</p>
OPTDUP_ENCRYPTION	<p>Specifies whether to encrypt the data during optimized duplication and replication. By default, files are not encrypted. If you want encryption, change the value to 1 on the MSDP storage server and on the MSDP load balancing servers. This parameter does not apply to clients.</p> <p>If you set this parameter to 1 on all hosts, the data is encrypted during transfer.</p> <p>Default value: OPTDUP_ENCRYPTION = 0</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See <a href="#">“About MSDP encryption”</a> on page 110.</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 Exchange Server files; <code>mdf</code> are SQL Server master database files, <code>ndf</code> are SQL Server 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: <code>file_extension:segment_size_in_KBs</code> pairs, separated by commas.</p> <p>See also <code>SEGKSIZE</code>.</p>

Table 5-31 pd.conf file parameters (continued)

Parameter	Description
PREFETCH_SIZE	<p>The size in bytes to use for the data buffer for restore operations.</p> <p>Default value: <code>PREFETCH_SIZE = 33554432</code></p> <p>Possible values: 0 to the computer's memory limit</p> <p><b>Note:</b> 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: <code>RESTORE_DECRYPT_LOCAL = 1</code></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: <code>SEGKSIZE = 128</code></p> <p>Possible values: 32 to 16384 KBs, increments of 32 only</p> <p><b>Warning:</b> 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 <code>PREFERRED_EXT_SEGKSIZE</code>.</p>

See [“About the MSDP pd.conf configuration file”](#) on page 169.

See [“Editing the MSDP pd.conf file”](#) on page 170.

## About the MSDP contentrouter.cfg file

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 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.

---

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

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

## About saving the MSDP 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 "spapasswd" " " 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 MSDP storage server configuration”](#) on page 184.

See [“Editing an MSDP storage server configuration file”](#) on page 184.

See [“Setting the MSDP storage server configuration”](#) on page 186.

## Saving the MSDP 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 MSDP storage server configuration”](#) on page 183.

See [“Recovering from an MSDP storage server disk failure”](#) on page 253.

See [“Recovering from an MSDP storage server failure”](#) on page 254.

### 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 an MSDP 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 5-32](#) shows the configuration lines that are required.



**Table 5-32** 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 <i>username</i> with the NetBackup Deduplication Engine user ID.
V7.0 "spapasswd" "password" string	Replace <i>password</i> with the password for the NetBackup Deduplication Engine user ID.

See [“About saving the MSDP storage server configuration”](#) on page 183.

See [“Recovering from an MSDP storage server disk failure”](#) on page 253.

See [“Recovering from an MSDP storage server failure”](#) on page 254.

**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 MSDP storage server configuration”](#) on page 184.

- 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 5-32](#)) 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 MSDP 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 an MSDP storage server disk failure”](#) on page 253.

See [“Recovering from an MSDP storage server failure”](#) on page 254.

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

See [“About saving the MSDP storage server configuration”](#) on page 183.

See [“Saving the MSDP storage server configuration”](#) on page 184.

See [“Editing an MSDP storage server configuration file”](#) on page 184.

---

**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 sshhostname -stype PureDisk -configlist file.txt`

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

For *sshhostname*, use the name of the storage server. For *file.txt*, use the name of the file that contains the configuration.

## About the MSDP 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.`, the configuration file name is `DedupeServer..cfg`.

The following is the location of the file:

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

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

## Deleting an MSDP 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 MSDP host configuration file”](#) on page 187.

**To delete the host configuration file**

- ◆ Delete the file on the deduplication host; 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..cfg`

## Resetting the MSDP registry

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

See [“Changing the MSDP storage server name or storage path”](#) on page 216.

---

**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 MSDP 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/openv/pdde/pdconfigure/cfg/userconfigs/pdregistry.cfg
   /etc/pdregistry.cfg
```

#### To reset the MSDP registry on Windows

- ◆ 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.

---

## About protecting the MSDP catalog

To increase availability, NetBackup provides a two-tier approach to protect the MSDP catalog, as follows:

Daily shadow copies	NetBackup automatically creates copies of the MSDP catalog. See <a href="#">“About the MSDP shadow catalog”</a> on page 188.
Catalog backup policy	Symantec provides a utility that you can use to configure a NetBackup policy that backs up the MSDP catalog. See <a href="#">“About the MSDP catalog backup policy”</a> on page 189.

See [“About recovering the MSDP catalog”](#) on page 250.

## About the MSDP shadow catalog

The NetBackup Deduplication Manager automatically creates a *shadow copy* of the catalog daily. The Deduplication Manager also builds a transaction log for each shadow copy. If NetBackup detects corruption in the MSDP catalog, the Deduplication Manager restores the catalog automatically from the most recent shadow copy. That restore process also plays the transaction log so that the recovered MSDP catalog is current.

By default, the NetBackup Deduplication Manager stores the shadow copies on the same volume as the catalog itself. Symantec recommends that you store the shadow copies on a different volume.

---

**Warning:** You can change the path only during initial MSDP configuration only. If you change it after MSDP backups exist, data loss may occur.

---

See [“Changing the MSDP shadow catalog path”](#) on page 191.

The NetBackup Deduplication Manager creates a shadow copy at 0340 hours daily, host time. To change the schedule, you must change the scheduler definition file.

See [“Changing the MSDP shadow catalog schedule”](#) on page 192.

By default, the NetBackup Deduplication Manager keeps five shadow copies of the catalog. You can change the number of copies.

See [“Changing the number of MSDP catalog shadow copies”](#) on page 193.

## About the MSDP catalog backup policy

Symantec recommends that you protect the MSDP catalog by backing it up. (A NetBackup catalog backup does not include the MSDP catalog.) The NetBackup Deduplication Catalog Policy Administration and Catalog Disaster Recovery utility (the `drcontrol` utility) configures a backup policy for the MSDP catalog. The policy also includes other important MSDP configuration information.

The MSDP catalog backups provide the second tier of catalog protection. The catalog backups are available if the shadow copies are not available or corrupt.

The following are the attributes for the catalog backup policy that the `drcontrol` utility creates:

Schedule	Weekly <b>Full Backup</b> and daily <b>Differential Incremental Backup</b> .
Backup window	6:00 A.M. to 6:00 P.M.
Retention	2 weeks

Backup selection    The following are the default catalog paths.

UNIX:

```
/database_path/databases/catalogshadow  
/storage_path/etc  
/storage_path/var  
/usr/opensv/lib/ost-plugins/pd.conf  
/usr/opensv/lib/ost-plugins/mtstrm.conf  
/database_path/databases/spa  
/database_path/databases/datacheck
```

Windows:

```
database_path\databases\catalogshadow  
storage_path\etc  
storage_path\var  
install_path\Veritas\NetBackup\bin\ost-plugins\pd.conf  
install_path\Veritas\NetBackup\bin\ost-plugins\mtstrm.conf  
database_path\databases\spa  
database_path\databases\datacheck
```

By default, NetBackup uses the same path for the storage and the catalog; the *database\_path* and the *storage\_path* are the same. If you configure a separate path for the deduplication database, the paths are different. Regardless, the `drcontrol` utility captures the correct paths for the catalog backup selections.

You should consider the following items carefully before you configure an MSDP catalog backup:

- Do not use the **Media Server Deduplication Pool** as the destination for the catalog backups. Recovery of the MSDP catalog from its **Media Server Deduplication Pool** is impossible.
- Use a storage unit that is attached to a NetBackup host other than the MSDP storage server.
- Use a separate MSDP catalog backup policy for each MSDP storage server. The `drcontrol` utility does not verify that the backup selections are the same for multiple storage servers. If the backup policy includes more than one MSDP storage server, the backup selection is the union of the backup selections for each host.
- You cannot use one policy to protect MSDP storage servers on both UNIX hosts and Windows hosts.  
UNIX MSDP storage servers require a Standard backup policy and Windows MSDP storage servers require an MS-Windows policy.

See [“Configuring an MSDP catalog backup”](#) on page 194.

See [“Updating an MSDP catalog backup policy”](#) on page 197.

## Changing the MSDP shadow catalog path

You can change the location of the catalog shadow copies. Symantec recommends that you store the copies on a different volume than both the *storage\_path* and the *database\_path*. (If you configured a separate path for the deduplication database, the paths are different.)

NetBackup stores the MSDP catalog shadow copies in the following location:

UNIX: */database\_path/databases/catalogshadow*

Windows: *database\_path\databases\catalogshadow*

---

**Warning:** You can change the shadow catalog path during initial MSDP configuration only. If you change it after MSDP backups exist, data loss may occur.

---

See [“About protecting the MSDP catalog”](#) on page 188.

### To change the MSDP catalog shadow path

- 1 Open the following file in a text editor:

UNIX: */storage\_path/etc/puredisk/spa.cfg*

Windows: *storage\_path\etc\puredisk\spa.cfg*

- 2 Find the `CatalogShadowPath` parameter and change the value to the wanted path.

The volume must be mounted and available.

- 3 After your changes, save the file.
- 4 Restart the NetBackup Deduplication Manager (*spad*).

- 5 Create the shadow catalog directories by invoking the following command on the MSDP storage server:

UNIX: `/usr/opensv/pdde/pdcr/bin/cacontrol --catalog backup all`

Windows: `install_path\Veritas\pdde\cacontrol --catalog backup all`

- 6 If an MSDP catalog backup policy exists, update the policy with the new shadow catalog directories. To do so, invoke the following command on the MSDP storage server:

UNIX: `/usr/opensv/pdde/pdcr/bin/drcontrol --update_policy --policy policy_name`

Windows: `install_path\Veritas\pdde\drcontrol --update_policy --policy policy_name`

## Changing the MSDP shadow catalog schedule

NetBackup automatically creates a copy of the MSDP catalog at 0340 hours daily, host time. You can change the default schedule.

See [“About protecting the MSDP catalog”](#) on page 188.



See “About protecting the MSDP catalog” on page 188.

**To change the number of MSDP catalog shadow copies**

- 1 Open the following file in a text editor:  
UNIX: `/storage_path/etc/puredisk/spa.cfg`  
Windows: `storage_path\puredisketc\spa.cfg`
- 2 Find the `CatalogBackupVersions` parameter and change the value to the wanted number of shadow copies. The valid values are 1 to 256, inclusive.
- 3 After your changes, save the file.
- 4 Restart the NetBackup Deduplication Manager (`spad`).

## Configuring an MSDP catalog backup

Use the following procedure to configure a backup policy for the NetBackup MSDP catalog.

See [“About protecting the MSDP data”](#) on page 37.

See [“Troubleshooting MSDP catalog backup”](#) on page 300.

### To configure an MSDP catalog backup

- 1 Verify that the MSDP storage server host (that is, the media server) is an additional server for the NetBackup master server. See **NetBackup Management > Host Properties > *masterserver\_name* > Servers > Additional Servers** in the NetBackup Administration Console.

If the storage server is not in the **Additional Servers** list, add the MSDP storage server host to the **Additional Servers** list. The host *must* be in the **Additional Servers** list and *cannot* be in the **Media Servers** list.

- 2 On the MSDP storage server, invoke the `drcontrol` utility and use the appropriate options for your needs. The following is the syntax for the utility:

UNIX: `/usr/opensv/pdde/pdcr/bin/drcontrol --new_policy --residence residence [--policy policy_name] [--client host_name] [--hardware machine_type] [--OS operating_system]`

Windows: `install_path\Veritas\pdde\drcontrol --new_policy --residence residence [--policy policy_name] [--client host_name] [--hardware machine_type] [--OS operating_system] [--NB_install_dir install_directory]`

Descriptions of the options are available in another topic.

See [“MSDP drcontrol options”](#) on page 195.

The utility creates a log file and displays its path in the command output.

See [“About MSDP log files”](#) on page 279.

## MSDP drcontrol options

The `drcontrol` utility resides in the following directories, depending on host type:

- UNIX: `/usr/opensv/pdde/pdcr/bin`
- Windows: `install_path\Veritas\pdde`

The `drcontrol` utility creates a log file.

See [“About MSDP log files”](#) on page 279.

[Table 5-33](#) describes the options for creating and updating an MSDP catalog backup policy.

**Table 5-33** MSDP `drcontrol` options for catalog backup and recovery

Option	Description
<code>--client host_name</code>	The client to back up (that is, the host name of the MSDP storage server).  Default: the value that <code>bpgetconfig CLIENT_NAME</code> returns.
<code>--hardware machine_type</code>	The hardware type or the computer type for the host.  Spaces are not allowed. If the string contains special characters, enclose it in double quotation marks (").  Default: Unknown.
	The pathname for the log file that the <code>drcontrol</code> utility creates. By default, the utility writes log files to <code>/storage_path/log/drcontrol/</code> .
<code>--NB_install_dir install_directory</code>	Windows only. Required option if NetBackup was installed in a location other than the default (C:\Program Files\Veritas).  If the string contains spaces or special characters, enclose it in double quotation marks ("). Do not use a trailing backslash in the <code>install_directory</code> string.
<code>--new_policy</code>	Create a new policy to protect the deduplication catalog on this host. If a policy with the given name exists already, the command fails.
<code>--OS operating_system</code>	The operating system for the host.  Spaces are not allowed. If the string contains special characters, enclose it in double quotation marks (").  Default: UNIX/Linux or MS-Windows.
<code>--policy policy_name</code>	The name for the backup policy.  Required with <code>--auto_recover_DR</code> and <code>--update_policy</code> ; optional with <code>--new_policy</code> .  Default: <code>Dedupe_Catalog_shorthostname</code>
<code>--refresh_shadow_catalog</code>	Deletes all existing shadow catalog copies and creates a new catalog shadow copy.
<code>--verbose</code>	Echo all <code>drcontrol</code> log statements to stdout.

**Table 5-33** MSDP `drcontrol` options for catalog backup and recovery  
(continued)

Option	Description
<code>--residence residence</code>	<p>The name of the storage unit on which to store the MSDP catalog backups.</p> <p>Do not use the <b>Media Server Deduplication Pool</b> as the destination for the catalog backups. Recovery of the MSDP catalog from its <b>Media Server Deduplication Pool</b> is impossible.</p> <p>Symantec recommends that you use a storage unit that is attached to a NetBackup host other than the MSDP storage server.</p>
<code>--update_policy</code>	<p>Update a policy, as follows:</p> <ul style="list-style-type: none"> <li>■ If the client name (of this media server) is not in the policy's client list, add the client name to the policy's client list.</li> <li>■ If you specify the <code>--OS</code> or <code>--hardware</code> options, replace the values currently in the policy with the new values.</li> <li>■ Update the backup selection based on the locations of the MSDP storage directories and configuration files. Therefore, if you modify any of the following, you must use this option to update the catalog backup policy: <ul style="list-style-type: none"> <li>■ Any of the following values in the <code>spa.cfg</code> file (section:variable pairs): <ul style="list-style-type: none"> <li>■ <code>StorageDatabase:CatalogShadowPath</code></li> <li>■ <code>StorageDatabase:Path</code></li> <li>■ <code>Paths:Var</code></li> </ul> </li> <li>■ The <code>spa.cfg</code> or <code>contentrouter.cfg</code> locations in the <code>pdregistry.cfg</code> file.</li> </ul> </li> </ul> <p>This option fails if there is no policy with the given policy name. It also fails if the existing policy type is incompatible with the operating system of the host on which you run the command.</p> <p>This option requires the <code>--policy policy_name</code> option.</p>

See [“Configuring an MSDP catalog backup”](#) on page 194.

## Updating an MSDP catalog backup policy

You can use any NetBackup method to update an MSDP catalog backup policy manually. However, you should use the NetBackup Deduplication Catalog Policy

Administration and Catalog Disaster Recovery(`drcontrol`) under the following circumstances:

- To add the client name of the storage server to the policy's client list.
- To update the `--os` value.
- To update the `--hardware` value.
- To update the backup selection if you modified any of the following configuration values:
  - Any of the following values in the `spa.cfg` file (section:variable pairs):
    - `StorageDatabase:CatalogShadowPath`
    - `StorageDatabase:Path`
    - `Paths:Var`
  - The `spa.cfg` or `contentrouter.cfg` locations in the `pdregistry.cfg` file.

See [“About protecting the MSDP data”](#) on page 37.

See [“Troubleshooting MSDP catalog backup”](#) on page 300.

#### To update an MSDP catalog backup

- ◆ On the MSDP storage server, invoke the `drcontrol` utility and use the appropriate options for your needs. The following is the syntax for an update operation:

**UNIX:** `/usr/opensv/pdde/pdcr/bin/drcontrol --update_policy --policy policy_name [--client host_name] [--hardware machine_type] [--OS operating_system]`

**Windows:** `install_path\Veritas\pdde\drcontrol --update_policy --policy policy_name [--client host_name] [--hardware machine_type] [--OS operating_system] [--OS operating_system] [--NB_install_dir install_directory]`

Descriptions of the options are available in another topic.

See [“MSDP drcontrol options”](#) on page 195.

The utility creates a log file and displays its path in the command output.

See [“About MSDP log files”](#) on page 279.

# Monitoring deduplication activity

This chapter includes the following topics:

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

## Monitoring the MSDP 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 MSDP deduplication rate:


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

On UNIX and Linux, you can use the NetBackup `bpdjobs` command to display the deduplication rate. However, you must configure it to do so.

See “To configure the `bpdbjobs` command to display the MSDP deduplication rate” on page 200.

To view the global MSDP 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 MSDP 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.

To configure the `bpdbjobs` command to display the MSDP deduplication rate

- ◆ Add a `DEDUPRATIO BPDBJOBS_COLDEFS` entry in the `bp.conf` file on the media server on which you run the command.

The `bpdbjobs` command then shows the deduplication rate in its output.



Many factors affect deduplication performance.

See [“About MSDP performance”](#) on page 31.

## Viewing MSDP job details

Use the NetBackup Administration Console to view deduplication job details.

To view MSDP 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 [“MSDP job details”](#) on page 201.

## MSDP job details

The **NetBackup Administration Console Job Details** dialog box shows the details of a deduplication job. The details depend on whether the job is media server deduplication or client-side deduplication job.

### Media server deduplication job details

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 MSDP\_Server. deduplicated the data (the **dedup** field shows the deduplication rate):

```
03/05/2013 10:37:06 - Info MSDP_Server. (pid=17498)
StorageServer=PureDisk:MSDP_Server.; Report=PDDO Stats
(multi-threaded stream used) for (MSDP_Server.):
scanned: 8930 KB, CR sent: 4363 KB, CR sent over FC: 0 KB, dedup: 51.1%,
cache hits: 36 (50.7%), rebased: 35 (49.3%)
```

### Client-side deduplication job details

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 **1/8/2013 11:58:09 PM** entry is for the client data; the **1/8/2013 11:58:19 PM** entry is for the metadata.

```
1/8/2013 11:54:21 PM - Info MSDP_Server.(pid=2220)
    Using OpenStorage client direct to backup from client
    Client_B. to MSDP_Server.
1/8/2013 11:58:09 PM - Info MSDP_Server.(pid=2220)
    StorageServer=PureDisk:MSDP_Server.; Report=PDDO
    Stats for (MSDP_Server.: scanned: 3423425 KB,
    CR sent: 122280 KB, dedup: 96.4%, cache hits: 49672 (98.2%)
1/8/2013 11:58:09 PM - Info MSDP_Server.(pid=2220)
    Using the media server to write NBU data for backup
    Client_B_1254987197. to MSDP_Server.
1/8/2013 11:58:19 PM - Info MSDP_Server.(pid=2220)
    StorageServer=PureDisk:MSDP_Server.; Report=PDDO
    Stats for (MSDP_Server.: scanned: 17161 KB,
    CR sent: 17170 KB, dedup: 0.0%, cache hits: 0 (0.0%)
```

## Field descriptions

[Table 6-1](#) describes the deduplication activity fields.

**Table 6-1** MSDP activity field descriptions

Field	Description
cache hits	<p>The percentage of data segments in the backup that is represented in the local fingerprint cache. The deduplication plug-in did not have to query the database about those segments</p> <p>If the <code>pd.conf</code> file <code>FP_CACHE_LOCAL</code> parameter is set to 0 on the storage, the <b>cache hits</b> output is not included for the jobs that run on the storage server.</p> <p>See <a href="#">"MSDP pd.conf file parameters"</a> on page 171.</p>
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"> <li>■ On a NetBackup client that deduplicates its own data (client-side deduplication).</li> <li>■ On a fingerprinting media server that deduplicates the data. The deduplication plug-in on the fingerprinting server sends the data to the storage server, which writes it to a <b>Media Server Deduplication Pool</b>.</li> <li>■ On a media server that then sends the data to a NetBackup 5000 series appliance for storage. (In NetBackup, a <b>PureDisk Storage Pool</b> represents the storage on a NetBackup 5000 series appliance.)</li> </ul>

Table 6-1 MSDP activity field descriptions (*continued*)

Field	Description
<b>CR sent over FC</b>	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.
<b>dedup</b>	The percentage of data that was stored already. That data is not stored again.
<b>multi-threaded stream used</b>	Indicates that the Deduplication Multi-Threaded Agent processed the backup. See <a href="#">“About the MSDP Deduplication Multi-Threaded Agent”</a> on page 52.
<b>PDDO stats</b>	Indicates that the job details are for storage on the following destinations: <ul style="list-style-type: none"><li>■ <b>Media Server Deduplication Pool</b></li><li>■ A NetBackup 5000 series appliance, exposed to NetBackup through a <b>PureDisk Deduplication Pool</b>.</li></ul>
<b>rebased</b>	The percentage of segments that were rebased (that is, defragmented) during the backup. Those segments had poor data locality.  NetBackup reports back up job completion only after backup rebasing is completed.  See <a href="#">“About MSDP storage rebasing”</a> on page 241.
<b>scanned</b>	The amount of data that the deduplication plug-in scanned.
<b>Using OpenStorage client direct to restore...</b>	Indicates that the restore travels over the client-direct data path and does not use NetBackup media server components to process the data.

## About MSDP 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 the 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 6-2](#) describes the options for monitoring capacity and usage.

**Table 6-2** Capacity and usage reporting

Option	Description
<b>Change Storage Server</b> dialog box	<p>The <b>Change Storage Server</b> dialog box <b>Properties</b> 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 <a href="#">“Monitoring the MSDP deduplication rate”</a> on page 199.</p>
<b>Disk Pools</b> window	<p>The <b>Disk Pools</b> window of the Administration Console displays the 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 <b>Change Storage Server</b> dialog box.</p> <p>To display the window, expand <b>Media and Device Management &gt; Devices &gt; Disk Pools</b>.</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 <a href="#">“About MSDP container files”</a> on page 205.</p> <p>See <a href="#">“Viewing storage usage within MSDP container files”</a> on page 205.</p>
<b>Disk Pool Status</b> report	<p>The <b>Disk Pool Status</b> report displays the state of the disk pool and usage information.</p> <p>See <a href="#">“Viewing MSDP disk reports”</a> on page 207.</p>

Table 6-2 Capacity and usage reporting (*continued*)

Option	Description
Disk Logs report	<p>The <b>Disk Logs</b> 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 <b>Disk Logs</b> report:</p> <p>The usage of one or more system resources has exceeded a warning level.</p> <p>By default, the threshold (high-water mark) for this message is at 98% capacity. No more data can be stored.</p> <p>See “<a href="#">Viewing MSDP disk reports</a>” on page 207.</p> <p>See “<a href="#">MSDP event codes and messages</a>” on page 301.</p>
The <code>nbdevquery</code> command	<p>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.</p> <p>See “<a href="#">Determining the MSDP disk volume state</a>” on page 232.</p>
NetBackup OpsCenter	<p>The NetBackup OpsCenter also provides information about storage capacity and usage.</p> <p>See <i>NetBackup OpsCenter Administrator's Guide</i>: <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>

## About MSDP 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.

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 storage usage within MSDP container files

The NetBackup `crcontrol` command reports on storage usage within containers.

**To view storage usage within MSDP 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:

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

The following is an example of the output:

```
***** Data Store statistics *****
Data storage      Raw    Size   Used   Avail  Use%
                  1.0T  988.9G 666.0G 322.9G 68%

Number of containers      : 2981
Average container size    : 219740494 bytes (209.56MB)
Space allocated for containers : 655046415189 bytes (610.06GB)
Reserved space           : 45360705536 bytes (42.25GB)
Reserved space percentage : 4.1%
```

From the command output, you can determine the following:

Raw	The raw size of the storage.
Size	<p>The size of the storage that NetBackup can use: the Raw size of the storage minus the file system Reserved space.</p> <p>If the file system has a concept of root reserved space (such as EXT3 or VxFS), that space cannot be used for storage. The <code>crcontrol</code> command does not include reserved space in the available space. Unlike the <code>crcontrol</code> command, some operating system tools report root reserved space as usable space.</p>
Used	The amount of deduplicated data that is stored on the file system. NetBackup obtains the file system used space from the operating system.
Avail	The Size minus the Used space.
Use%	The Used space divided by the Size.

## Viewing MSDP 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 6-3](#) describes the disk reports available.

**Table 6-3** 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>
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 <a href="#">“About MSDP data integrity checking”</a> on page 236.</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 and Symantec Backup Exec are Symantec applications that use deduplication.</p>
Disk Storage Unit Status	<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 <b>Media and Device Management &gt; Devices &gt; Disk Pools</b>.</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 <b>Disk Pool Status</b> 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**.

## About monitoring MSDP processes

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

See “[MSDP server components](#)” on page 263.

**Table 6-4** Where to monitor the main MSDP processes

What	Where to monitor it
NetBackup Deduplication Engine	On Windows systems, in the NetBackup Administration Console <b>Activity Monitor Services</b> tab.  On UNIX, the NetBackup Deduplication Engine appears as <code>spoold</code> in the Administration Console <b>Activity Monitor Daemons</b> tab.  The NetBackup <code>bpps</code> command also shows the <code>spoold</code> process.
NetBackup Deduplication Manager	On Windows systems, NetBackup Deduplication Manager in the <b>Activity Monitor Services</b> tab.  On UNIX, the NetBackup Deduplication Manager appears as <code>spad</code> in the Administration Console <b>Activity Monitor Daemons</b> tab.  The NetBackup <code>bpps</code> command also shows the <code>spad</code> process.

## 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 6-5 Auto Image Replication jobs in the Activity Monitor

Job type	Description
<b>Replication</b>	<p>The job that replicates a backup image to a target master displays in the Activity Monitor as a <b>Replication</b> job. The <b>Target Master</b> label displays in the <b>Storage Unit</b> column for this type of job.</p> <p>Similar to other <b>Replication</b> 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>
<b>Import</b>	<p>The job that imports a backup copy into the target master domain displays in the Activity Monitor as an <b>Import</b> job. An <b>Import</b> job can import multiple copies in one instance. The detailed status for an <b>Import</b> job contains a list of processed backup IDs and a list of failed backup IDs.</p> <p><b>Note:</b> If the master servers in the source and target domains are not at the same NetBackup version, the following error can occur under certain circumstances: Failed to auto create data classification.</p> <p>This error occurs if the master server in the source domain is at a NetBackup version earlier than 7.6 and the data classification of <b>Any</b> is used. If the master server in the target domain is at NetBackup 7.6, use a different data classification in the source domain or the Import job fails.</p> <p>Note that a successful replication does not confirm that the image was imported at the target master.</p> <p>If the data classifications are not the same in both domains, the <b>Import</b> job fails and NetBackup does not attempt to import the image again.</p> <p>Failed <b>Import</b> jobs fail with a status 191 and appear in the <b>Problems</b> report when run on the target master server.</p> <p>The image is expired and deleted during an <b>Image Cleanup</b> job. Note that the originating domain (Domain 1) does not track failed imports.</p>

# Managing deduplication

This chapter includes the following topics:

- [Managing MSDP servers](#)
- [Managing NetBackup Deduplication Engine credentials](#)
- [Managing Media Server Deduplication Pools](#)
- [Deleting backup images](#)
- [About MSDP queue processing](#)
- [Processing the MSDP transaction queue manually](#)
- [About MSDP data integrity checking](#)
- [Configuring MSDP data integrity checking behavior](#)
- [About managing MSDP storage read performance](#)
- [About MSDP storage rebasing](#)
- [About the MSDP data removal process](#)
- [Resizing the MSDP storage partition](#)
- [How MSDP restores work](#)
- [Configuring MSDP restores directly to a client](#)
- [About restoring files at a remote site](#)
- [About restoring from a backup at a target master domain](#)
- [Specifying the restore server](#)

## Managing MSDP servers

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

See [“Viewing MSDP storage servers”](#) on page 211.

See [“Determining the MSDP storage server state”](#) on page 211.

See [“Viewing MSDP storage server attributes”](#) on page 212.

See [“Setting MSDP storage server attributes”](#) on page 213.

See [“Changing MSDP storage server properties”](#) on page 214.

See [“Clearing MSDP storage server attributes”](#) on page 215.

See [“About changing the MSDP storage server name or storage path”](#) on page 216.

See [“Changing the MSDP storage server name or storage path”](#) on page 216.

See [“Removing an MSDP load balancing server”](#) on page 218.

See [“Deleting an MSDP storage server”](#) on page 219.

See [“Deleting the MSDP storage server configuration”](#) on page 220.

## Viewing MSDP storage servers

Use the NetBackup Administration Console to view a list of deduplication storage servers already configured.

### To view MSDP 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 **Server Type** column.

## Determining the MSDP 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 MSDP 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 MSDP 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 MSDP 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 MSDP 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 a Media Server Deduplication Pool attribute”](#) on page 225.

### To set a MSDP 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, <b>OptimizedImage</b> 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 MSDP storage server attributes”](#) on page 212.

See [“About MSDP optimized synthetic backups”](#) on page 29.

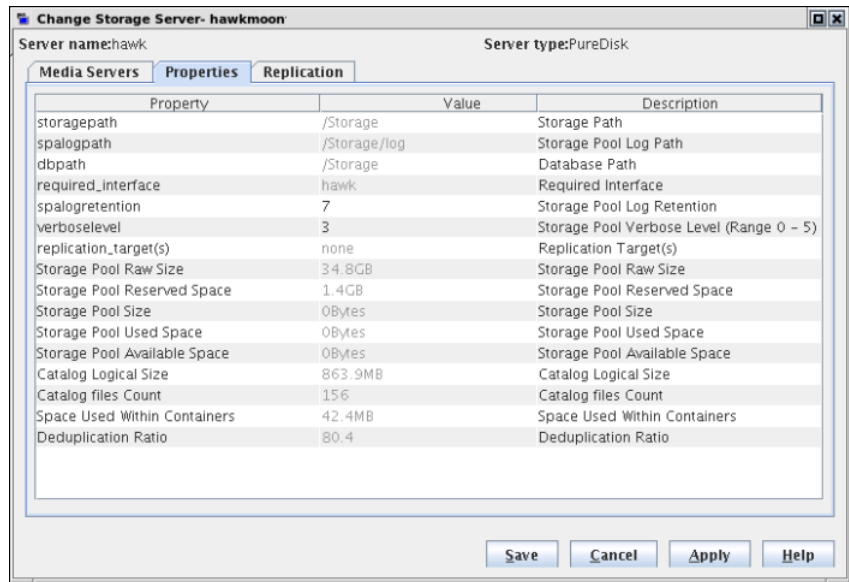
## Changing MSDP storage server properties

You can change the retention period and logging level for the NetBackup Deduplication Manager.

To change MSDP 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 MSDP storage server attributes

Use the `nbdevconfig` command to remove storage server attributes.

### To clear MSDP 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 the  
`attribute`            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 MSDP 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 MSDP storage server name or storage path”](#) on page 216.
- You want to change the storage path. To do so, you must reconfigure the storage server with the new path.  
See [“Changing the MSDP storage server name or storage path”](#) on page 216.
- 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.  
See [“Recovering from an MSDP storage server failure”](#) on page 254.

## Changing the MSDP 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 MSDP storage server name or storage path”](#) on page 216.



**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, Volume I</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>
Step 2	Expire the backup images	Expire all backup images that reside on the deduplication disk storage.  <b>Warning:</b> 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, Volume I</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>
Step 3	Delete the storage units that use the disk pool	See the <i>NetBackup Administrator's Guide, Volume I</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>
Step 4	Delete the disk pool	See "Deleting a Media Server Deduplication Pool" on page 234.
Step 5	Delete the deduplication storage server	See "Deleting an MSDP storage server" on page 219.
Step 6	Delete the configuration	Delete the deduplication configuration.  See "Deleting the MSDP storage server configuration" on page 220.
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 an MSDP host configuration file" on page 187.
Step 8	Delete the identity file and the file system table file	Delete the following files from the MSDP storage server, depending on operating system:  UNIX:  <code>/storage_path/data/.identify</code> <code>/storage_path/etc/puredisk/fstab.cfg</code>  Windows:  <code>storage_path\data\.identify</code> <code>storage_path\etc\puredisk\fstab.cfg</code>

**Table 7-1** Changing the storage server name or storage path (*continued*)

Step	Task	Procedure
Step 9	Change the storage server name or the storage location	See the computer or the storage vendor's documentation. See <a href="#">"Use fully qualified domain names"</a> on page 33. See <a href="#">"MSDP storage path properties"</a> on page 78.
Step 10	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. See <a href="#">"Configuring MSDP server-side deduplication"</a> on page 49.
Step 11	Import the backup images	See the <i>NetBackup Administrator's Guide, Volume I</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>

## Removing an MSDP load balancing server

You can remove a load balancing server from a deduplication node. The media server no longer deduplicates client data.

See ["About MSDP storage servers"](#) on page 21.

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 *NetBackup Administrator's Guide, Volume II*:

<http://www.symantec.com/docs/DOC5332>

### To remove a media server from a MSDP 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 an MSDP 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 an MSDP 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 MSDP storage server name or storage path”](#) on page 216.

## Deleting the MSDP storage server configuration

Use this procedure to delete a deduplication storage server configuration. The script that is used in this procedure deletes the active configuration and returns the configuration files to their installed, preconfigured 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 MSDP storage server name or storage path”](#) on page 216.

See [“Deactivating MSDP”](#) on page 261.

#### To delete the MSDP storage server configuration

- 1 Use the NetBackup Administration Console to stop the NetBackup Deduplication Engine (`spoold`) and the NetBackup Deduplication Manager (`spad`).
- 2 On the storage server, run one of the following scripts, depending on your operating system:

UNIX:

```
/usr/openv/pdde/pdconfigure/scripts/installers/PDDE_deleteConfig.sh
```

Windows: `install_path\Program`

```
Files\Veritas\pdde\PDDE_deleteConfig.bat
```

The command output includes the following:

```
**** Starting PDDE_deleteConfig.sh ****
You need to stop the spad and spoold daemons to proceed
This script will delete the PDDE configuration on this system
Would you want to continue? [ y | n ]
```

- 3 Type **y** and then press Enter.

## About shared memory on Windows MSDP 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 221.

See [“Adding NetBackup Deduplication Engine credentials”](#) on page 222.

See [“Changing NetBackup Deduplication Engine credentials”](#) on page 222.

See [“Deleting credentials from a load balancing server”](#) on page 222.

## 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 an MSDP load balancing server”](#) on page 168.

**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:

On Windows:

```
install_path\Veritas\NetBackup\Volmgr\bin\tpconfig -add  
-storage_server sshostname -stype PureDisk -sts_user_id UserID  
-password PassWord
```

On UNIX/Linux:

```
/usr/opensv/volmgr/bin/tpconfig -add -storage_server sshostname  
-stype PureDisk -sts_user_id UserID -password PassWord
```

For *sshostname*, 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 the NetBackup Deduplication Engine credentials”](#) on page 28.

## 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 an MSDP load balancing server”](#) on page 218.

### To delete credentials from a load balancing server

- ◆ On the load balancing server, run the following command:

On Windows:

```
install_path\Veritas\NetBackup\Volmgr\bin\tpconfig -delete
-storage_server sshostname -stype PureDisk -sts_user_id UserID
```

On UNIX/Linux:

```
/usr/openv/volmgr/bin/tpconfig -delete -storage_server sshostname
-stype PureDisk -sts_user_id UserID
```

For *sshostname*, use the name of the storage server.

## Managing Media Server Deduplication Pools

After you configure NetBackup deduplication, you can perform various tasks to manage your deduplication disk pools.

- See [“Viewing Media Server Deduplication Pools”](#) on page 223.
- See [“Changing a Media Server Deduplication Pool properties”](#) on page 226.
- See [“Determining the Media Server Deduplication Pool state”](#) on page 224.
- See [“Changing a Media Server Deduplication Pool state”](#) on page 224.
- See [“Determining the MSDP disk volume state”](#) on page 232.
- See [“Changing the MSDP disk volume state”](#) on page 232.
- See [“Viewing Media Server Deduplication Pool attributes”](#) on page 224.
- See [“Setting a Media Server Deduplication Pool attribute”](#) on page 225.
- See [“Clearing a Media Server Deduplication Pool attribute”](#) on page 231.
- See [“Resizing the MSDP storage partition”](#) on page 245.
- See [“Inventorying a NetBackup disk pool”](#) on page 233.
- See [“Deleting a Media Server Deduplication Pool”](#) on page 234.

## Viewing Media Server Deduplication 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 Media Server Deduplication 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 a Media Server Deduplication 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 MSDP 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 disk pools for NetBackup deduplication”](#) on page 91.

See [“Determining the Media Server Deduplication Pool state”](#) on page 224.

See [“Media Server Deduplication Pool properties”](#) on page 100.

See [“Configuring a disk pool for deduplication”](#) on page 92.

## Viewing Media Server Deduplication Pool attributes

Use the NetBackup `nbdevquery` command to view deduplication pool attributes.



### To view MSDP 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. (UP)
```

This example output is shortened; more flags may appear in actual output.

## Setting a Media Server Deduplication Pool attribute

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 MSDP storage server attributes”](#) on page 213.

### To set a MSDP 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:

`-changedp`                      The name of the disk pool.  
*pool\_name*

`-setattribute`                  The *attribute* is the name of the argument that represents the new functionality.  
*attribute*

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 Media Server Deduplication Pool attributes”](#) on page 224.

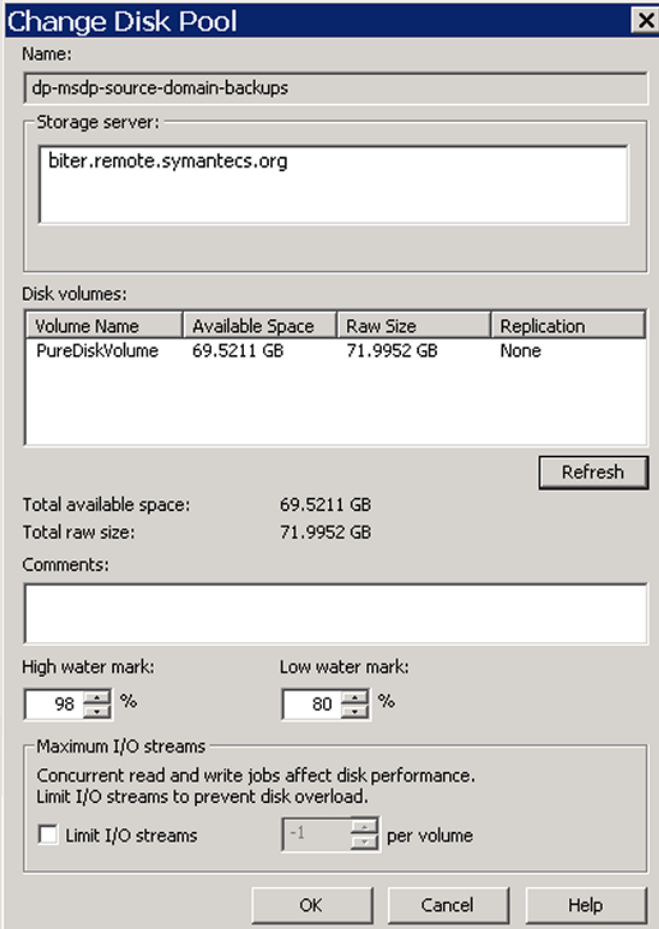
## Changing a Media Server Deduplication Pool properties

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**.



**Change Disk Pool**

Name:

Storage server:

Disk volumes:

Volume Name	Available Space	Raw Size	Replication
PureDiskVolume	69.5211 GB	71.9952 GB	None

Total available space: 69.5211 GB  
Total raw size: 71.9952 GB

Comments:

High water mark:  %  
Low water mark:  %

Maximum I/O streams  
Concurrent read and write jobs affect disk performance.  
Limit I/O streams to prevent disk overload.

☐ Limit I/O streams  per volume

- 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 228.

- 5 Change the other properties as necessary.

See [“Media Server Deduplication Pool properties”](#) on page 100.

- 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 queries the storage server for changes when you either click the **Refresh** button in the **Change Disk Pool** dialog box or when you configure a new disk pool for the storage server.

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, ask your storage administrator to revert the changes 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.

The following table describes the possible outcomes and 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 <b>Change Disk Pool</b> 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)

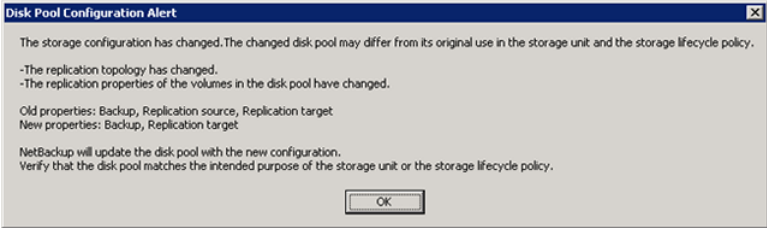
Outcome	Description
The replication properties of all of the volumes changed, but they are still consistent.	<p>A <b>Disk Pool Configuration Alert</b> pop-up box notifies you that the properties of all of the volumes in the disk pool changed, but they are all the same (homogeneous).</p> <div></div> <p>You must click <b>OK</b> in the alert box, after which the disk pool properties in the <b>Change Disk Pool</b> 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 <b>Change Disk Pool</b> dialog box. You can add those new volumes to the disk pool.</p> <p>In the <b>Change Disk Pool</b> dialog box, select one of the following two choices:</p> <ul style="list-style-type: none"><li>■ <b>OK.</b> To accept the disk pool changes, click <b>OK</b> in the <b>Change Disk Pool</b> 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.</li><li>■ <b>Cancel.</b> To discard the changes, click <b>Cancel</b> in the <b>Change Disk Pool</b> 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.</li></ul>

Table 7-2 Refresh outcomes (continued)

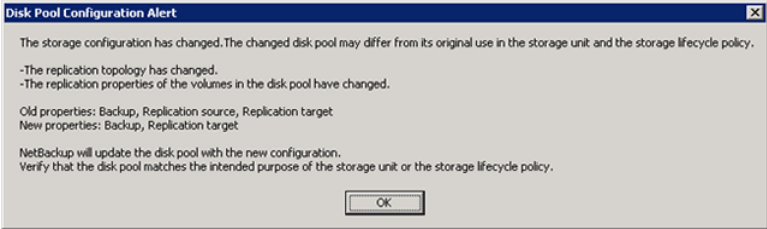

Outcome	Description
The replication properties of the volumes changed, and they are now inconsistent.	<p>A <b>Disk Pool Configuration Error</b> pop-up box notifies you that the replication properties of some of the volumes in the disk pool changed. The properties of the volumes in the disk pool are not homogeneous.</p> <div></div> <p>You must click <b>OK</b> in the alert box.</p> <p>In the <b>Change Disk Pool</b> 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 <a href="#">“Viewing the replication topology for Auto Image Replication”</a> on page 142.</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 <b>Change Disk Pool</b> dialog box, click <b>OK</b> or <b>Cancel</b> to exit the <b>Change Disk Pool</b> 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 <b>Disk Pool Configuration Alert</b> pop-up box notifies you that an existing volume or volumes was deleted from the storage device:</p> <div><div><div>Disk Pool Configuration Alert</div><div><div><div>An existing volume in this disk pool cannot be found on the storage device and is no longer available to NetBackup. The volume might be offline or deleted. If deleted, any data on that volume is lost.</div><div>Volume(s) deleted: dv02</div><div>Refer to documentation for information on how to resolve this issue.</div></div><div>OK</div></div></div></div> <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"><li>■ Use the <code>bpimmedia</code> command or the Images on Disk report to display the images on the specific volume.</li><li>■ Expire the images on the volume.</li><li>■ Use the <code>nbdevconfig</code> command to set the volume state to DOWN so NetBackup does not try to use it.</li></ul>

## Clearing a Media Server Deduplication Pool attribute

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> <code><i>pool_name</i></code>	The name of the disk pool.
<code>-setattribute</code> <code><i>attribute</i></code>	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/openv/netbackup/bin/admincmd`
- **Windows:** `install_path\NetBackup\bin\admincmd`

## Determining the MSDP disk volume state

Use the NetBackup `nbdevquery` command to determine the state of the volume in a deduplication disk pool. NetBackup exposes all of the storage for MSDP as a single volume, **PureDiskVolume**. The command shows the properties and attributes of the **PureDiskVolume**.

### To determine MSDP disk volume state

- ◆ Display the volume state by using the following command:

**UNIX:** `/usr/openv/netbackup/bin/admincmd/nbdevquery -listdv -stype PureDisk -U -dp disk_pool_name`

**Windows:** `install_path\NetBackup\bin\admincmd\nbdevquery -listdv -stype PureDisk -U -dp disk_pool_name`

The *state* is either UP or DOWN.

The following is example output

```
Disk Pool Name      : MSDP_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 MSDP disk volume state

The disk volume state is **UP** or **DOWN**. NetBackup exposes all of the storage for MSDP as a single volume, **PureDiskVolume**.



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 MSDP disk volume state

- ◆ Change the disk volume state; the following is the command syntax:

**UNIX:** `/usr/opensv/netbackup/bin/admincmd/nbdevconfig -changestate  
-stype PureDisk -dp disk_pool_name -dv PureDiskVolume -state state`

**Windows:** `install_path\NetBackup\bin\admincmd\nbdevconfig  
-changestate -stype PureDisk -dp disk_pool_name -dv PureDiskVolume  
-state state`

For the `-state`, specify either **UP** or **DOWN**.

## Inventorying a NetBackup disk pool

An inventory of a NetBackup disk pool reads the capacity of the disk volumes in the pool. An inventory operation lets you update NetBackup with the new capacity values if you do the following:

- Increase or decrease the size of the disk volumes in a disk pool.
- Add volumes to or remove volumes from a disk pool.

How you increase or decrease the underlying storage capacity depends on your storage implementation. You must complete that process before you inventory the disk pool.

#### To inventory a NetBackup disk pool

- 1 in the **NetBackup Administration Console**, select **Media and Device Management > Devices > Disk Pools**.
- 2 On the **Actions** menu, select **Inventory Disk Pools**.
- 3 In the **Inventory Disk Pool** dialog box, select the disk pool to inventory and then click **Start Inventory**.
- 4 To update the NetBackup catalog with the capacity values returned by the inventory, click **Update Configuration**.
- 5 To inventory another disk pool, go to step 3.
- 6 To exit, click **Close**.

## Deleting a Media Server Deduplication 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 [“Cannot delete an MSDP disk pool”](#) on page 293.

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 an MSDP 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 [“About the MSDP data removal process”](#) on page 244.

### 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.

## About MSDP 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. By default, queue processing occurs every 12 hours, 20 minutes past the hour.

Primarily, the transaction queue contains clean-up and integrity checking transactions. These transactions update the reference database.

Queue processing writes status information to the deduplication engine `stored.log` file.

See [“About MSDP log files”](#) on page 279.

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 MSDP transaction queue manually”](#) on page 235.

See [“About MSDP server requirements”](#) on page 22.

## Processing the MSDP transaction queue manually

NetBackup maintains a queue for MSDP database transactions.

See [“About MSDP queue processing”](#) on page 235.

Usually, you should not need to run the deduplication database transaction queue processes manually. However, when you recover the MSDP catalog from a backup, you must process the MSDP transaction queue. Processing the transaction queue is part of a larger process.

### To process the MSDP transaction queue manually

- 1 On the MSDP storage server, run the following command:

UNIX: `/usr/opensv/pdde/pdcr/bin/crcontrol --processqueue`

Windows: `install_path\Veritas\pdde\Crcontrol.exe --processqueue`

- 2 To determine if the queue processing is still active, run the following command:

UNIX: `/usr/opensv/pdde/pdcr/bin/crcontrol --processqueueinfo`

Windows: `install_path\Veritas\pdde\Crcontrol.exe --processqueueinfo`

If the output shows `Busy : yes`, the queue is still active.

- 3 To examine the results, run the following command (number 1 not lowercase letter l):

UNIX: `/usr/opensv/pdde/pdcr/bin/crcontrol --dsstat 1`

Windows: `install_path\Veritas\pdde\Crcontrol.exe --dsstat 1`

The command may run for a long time; if you omit the 1, results return more quickly but they are not as accurate.

## About MSDP 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 and actions:

- Automatically constrains data loss or corruption to ensure that new backups are intact.
- Automatically runs a cyclic redundancy check (CRC) for the data containers.
- Automatically collects and cleans up storage garbage.
- Automatically recovers the container-based reference database (or parts of the database) if it is corrupt or missing.
- Automatically finds storage leaks and fixes them.

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 MSDP disk reports”](#) on page 207.

The data integrity message code is 1057.

See [“MSDP event codes and messages”](#) on page 301.

NetBackup writes the integrity checking activity messages to the NetBackup Deduplication Engine `storaged.log` file.

See [“About MSDP log files”](#) on page 279.

You can configure some of the data integrity checking behaviors.

See [“Configuring MSDP data integrity checking behavior”](#) on page 237.

## Configuring MSDP data integrity checking behavior

NetBackup performs several data integrity checks. You can configure the behavior of the integrity checks.

Two methods exist to configure MSDP data integrity checking behavior, as follows:

- Run a command.  
See [“To configure data integrity checking behavior by using a command”](#) on page 237.
- Edit configuration file parameters.  
See [“To configure data integrity checking behavior by editing the configuration files”](#) on page 239.

---

**Warning:** Symantec recommends that you do not disable the data integrity checking. If you do so, NetBackup cannot find and repair or report data corruption.

---

See [“About MSDP data integrity checking”](#) on page 236.

See [“MSDP data integrity checking configuration parameters”](#) on page 239.

### 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:

Enable	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a enabledataintegritycheck</code></p> <p>Windows: <code>install_path\Veritas\pdde\pddecfg -a enabledataintegritycheck</code></p>
Disable	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a disabledataintegritycheck</code></p> <p>Windows: <code>install_path\Veritas\pdde\pddecfg -a disabledataintegritycheck</code></p>
Get the status	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a getdataintegritycheck</code></p> <p>Windows: <code>install_path\Veritas\pdde\pddecfg -a getdataintegritycheck</code></p>

- **Cyclic redundancy checking.** Use the following commands to configure behavior:

Enable	<p>CRC 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>
Enable fast checking	<p>Fast check CRC mode begins the check from container 64 and does not sleep between checking containers.</p> <p>When the fast CRC ends, CRC behavior reverts to the behavior before fast checking was invoked.</p> <p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckrestart</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --crccheckrestart</code></p>

```
Get the status  UNIX: /usr/openv/pdde/pdcr/bin/crcontrol
                --crccheckstate

                Windows: install_path\Veritas\pdde\Crcontrol.exe
                --crccheckstate
```

### To configure data integrity checking behavior by editing the configuration files

- 1 Use a text editor to open the `contentrouter.cfg` file or the `spa.cfg` file, which control the data integrity checking behavior.

The files reside in the following directories:

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

- 2 To change a parameter, specify a new value.

See “[MSDP data integrity checking configuration parameters](#)” on page 239.

- 3 Save and close the file.

- 4 Restart the NetBackup Deduplication Engine and the NetBackup Deduplication Manager.

Use the **NetBackup Administration Console Activity Monitor Services** tab (Windows) or **Daemons** tab (UNIX).

## MSDP data integrity checking configuration parameters

The configuration file parameters that control the deduplication data integrity checking are in two different configuration files, as follows:

- The `contentrouter.cfg` file.  
The parameters are described in [Table 7-3](#).  
See “[About the MSDP contentrouter.cfg file](#)” on page 182.
- The `spa.cfg` file.  
The parameters are described in [Table 7-3](#).

Those files reside in the following directories:

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

---

**Warning:** Symantec recommends that you do not disable the data integrity checking. If you do so, NetBackup cannot find and repair or report data corruption.

---

See [“About MSDP data integrity checking”](#) on page 236.

**Table 7-3** The contentrouter.cfg file parameters for data integrity checking

Setting	Default	Description
EnableCRCCheck	true	<p>Enable or disable cyclic redundancy checking (CRC) of the data container files.</p> <p>The possible values are <code>true</code> or <code>false</code>.</p> <p>CRC occurs only when no backup, restore, or queue processing jobs are running.</p>
CRCCheckSleepSeconds	5	<p>The time in seconds to sleep between checking containers.</p> <p>The longer the sleep interval, the more time it takes to check containers.</p>
CRCCheckBatchNum	40	<p>The number of containers to check each time.</p> <p>The greater the number of containers, the less time it takes to check all containers, but the more system resources it takes.</p>
ShutdownCRWhenError	false	<p>Stops the NetBackup Deduplication Manager when a data loss is discovered.</p> <p>This parameter is reserved for debugging purposes by Symantec Support Representatives.</p> <p>The possible values are <code>true</code> or <code>false</code>.</p>
GarbageCheckRemainDCCount	100	<p>The number of containers from failed jobs not to check for garbage. A failed backup or replication job still produces data containers. Because failed jobs are retried, retaining those containers means NetBackup does not have to send the fingerprint information again. As a result, retried jobs consume less time and fewer system resources than when first run.</p>

**Table 7-4** spa.cfg file parameters for data integrity checking

Setting	Default	Description
EnableDataCheck	true	<p>Enable or disable data consistency checking.</p> <p>The possible values are <code>True</code> or <code>False</code>.</p>
DataCheckDays	14	<p>The number of days to check the data for consistency.</p> <p>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.</p>



Table 7-4 spa.cfg file parameters for data integrity checking (*continued*)

Setting	Default	Description
EnableDataCheckAlert	true	Enable or disable alerts.  If <code>true</code> , NetBackup writes a message to the Disk Logs report when it detects a lost data segment.  See <a href="#">“About MSDP log files”</a> on page 279.

## About managing MSDP 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, duplication jobs, and replication jobs.

In most cases, you should change configuration file options only when directed to do so by a Symantec support representative.

### Defragment the storage

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.

If you upgrade from a NetBackup release earlier than 7.5, rebasing may affect your deduplication performance temporarily.

See [“About MSDP storage rebasing”](#) on page 241.

### Decrypt the data on the client rather than the server

The `RESTORE_DECRYPT_LOCAL` parameter in the `pd.conf` file specifies on which host to decrypt and decompress the data during restore operations.

See [“About the MSDP pd.conf configuration file”](#) on page 169.

See [“MSDP pd.conf file parameters”](#) on page 171.

## About MSDP storage rebasing

During an initial backup, NetBackup writes the data segments from a backup to as few container files as possible. Read performance is best when the data segments from a client backup are close to each other on deduplication storage. NetBackup consumes less time finding and reassembling backed up files when their segments are near each other.

However, the data segments in a backup may become scattered across the disk storage each time the client is backed up. Such scattering is a normal consequence of deduplication.

NetBackup includes a process, called *rebasing*, that helps to maintain the data segments in as few container files as possible. Rebasing improves performance for the operations that read from the storage, such as restores and duplications. NetBackup writes all of the data segments from a backup into new container files even though the segments exist on storage already. Future backups then refer to the new copies of those segments rather than the old copies until any changes because of future rebasing. Deduplication rates for the backup jobs that perform rebasing are lower than for the jobs that do not rebase the data.

NetBackup reports back up job completion after rebasing is completed.

After the rebasing, NetBackup reclaims the storage space that the rebased data segments used.

[Table 7-5](#) describes the rebasing operations.

**Table 7-5**           Types of rebasing

Type	Description
Normal backup rebasing	<p>The rebasing that occurs during a backup if the normal rebasing criteria are met, as follows:</p> <ul style="list-style-type: none"><li>■ The container has been rebased within the last three months.</li><li>■ For that backup, the data segments in the container consume less space than the <code>FP_CACHE_PERIOD_REBASING_THRESHOLD</code> value. The <code>FP_CACHE_PERIOD_REBASING_THRESHOLD</code> parameter is in the <code>pd.conf</code> file.</li></ul> <p>See <a href="#">“MSDP pd.conf file parameters”</a> on page 171.</p> <p>Backup rebasing occurs only for the full backups that pass through the normal MSDP backup process. For example, the NetBackup Accelerator backups do not pass through the MSDP backup process.</p>

**Table 7-5** Types of rebasing (*continued*)

Type	Description
Periodic backup rebasing	<p>The rebasing that occurs during a backup if the periodic rebasing criteria are met, as follows:</p> <ul style="list-style-type: none"> <li>■ The container has not been rebased within the last three months.</li> <li>■ For that backup, the data segments in the container consume less space than the <code>FP_CACHE_REBASING_THRESHOLD</code> value. The <code>FP_CACHE_REBASING_THRESHOLD</code> parameter is in the <code>pd.conf</code> file.</li> </ul> <p>See “<a href="#">MSDP pd.conf file parameters</a>” on page 171.</p> <p>Backup rebasing occurs only for the full backups that pass through the normal MSDP backup process. For example, the NetBackup Accelerator backups do not pass through the MSDP backup process.</p>
Server-side rebasing	<p>The storage rebasing that occurs on the server if the rebasing criteria are met. Server-side rebasing includes the deduplicated data that does not pass through the normal MSDP backup process. For example, the NetBackup Accelerator backups do not pass through the MSDP backup process.</p> <p>The following parameters in the <code>contentrouter.cfg</code> file control the server-side rebasing behavior:</p> <ul style="list-style-type: none"> <li>■ <code>RebaseScatterThreshold</code></li> <li>■ <code>RebaseMinContainers</code></li> <li>■ <code>RebaseMaxPercentage</code></li> <li>■ <code>RebaseMaxTime</code></li> </ul> <p>See “<a href="#">MSDP server-side rebasing parameters</a>” on page 243.</p>

## MSDP server-side rebasing parameters

[Table 7-6](#) describes the parameters that control server-side rebasing.

See “[About MSDP storage rebasing](#)” on page 241.

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

The parameters are stored in the `contentrouter.cfg` file.

See “[About the MSDP contentrouter.cfg file](#)” on page 182.

**Table 7-6** The server-side rebasing parameters

Parameter	Description
<code>RebaseMaxPercentage</code>	<p>The maximum percentage of the data segments to be rebased in a file. For any file, if the percentage of the data segments reaches this threshold, the remainder of the data segments are not rebased.</p> <p>By default, this parameter is <code>RebaseMaxPercentage=5</code>.</p>
<code>RebaseMaxTime</code>	<p>The maximum time span in seconds of data segments to be rebased in a file. If this threshold is reached, NetBackup does not rebase the remainder of the data segments.</p> <p>By default, this parameter is <code>RebaseMaxTime=150</code>.</p>
<code>RebaseMinContainers</code>	<p>The minimum number of containers in which a file's data segments are stored for the file to be eligible for rebasing. If the number of containers in which a file's data segments are stored is less than <code>RebaseMinContainers</code>, NetBackup does not rebase the data segments.</p> <p>By default, this parameter is <code>RebaseMinContainers=4</code>.</p>
<code>RebaseScatterThreshold</code>	<p>The data locality threshold for a container. If the total size of a file's data segments in a container is less than <code>RebaseScatterThreshold</code>, NetBackup rebases all of the file's data segments.</p> <p>By default, this parameter is <code>RebaseScatterThreshold=64MB</code>.</p>

## About the MSDP data removal process

The data removal process removes the data segments that comprise a NetBackup backup image. Only those segments that are not referred to by a backup image are removed.

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 in the deduplication catalog and adds a removal request to the NetBackup Deduplication Engine's transaction queue.  
From this point on, the expired backup image is no longer accessible.
- When the NetBackup Deduplication Engine processes the queue, all of the removal requests are processed. A removal request for the image is not generated again.  
During the queue processing, the Deduplication Engine reclaims some of the storage space on which the data segments reside. Some is reclaimed during

data compaction. If a different backup image requires a data segment, the segment is not removed.

Various internal parameters control whether a container file is compacted.

See “[About MSDP container files](#)” on page 205.

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. Garbage collection occurs during data integrity checking.

See “[About MSDP data integrity checking](#)” on page 236.

See “[Deleting backup images](#)” on page 234.

## Resizing the MSDP 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 MSDP 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, Volume I*:

<http://www.symantec.com/docs/DOC5332>

- 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, Volume I*:

<http://www.symantec.com/docs/DOC5332>

- 7 Restart the deduplication jobs.

See “[About adding additional MSDP storage](#)” on page 44.

# How MSDP restores work

The following two methods exist to for MSDP restore operations:

**Table 7-7** MSDP restore types

Type	Description
Normal restore	<p>The MSDP storage server first <i>rehydrates</i> (that is, reassembles) the data. NetBackup then chooses the least busy media server to move the data to the client. (NetBackup chooses the least busy media server from those that have credentials for the NetBackup Deduplication Engine.) The media server <code>bptm</code> process moves the data to the client.</p> <p>The following media servers have credentials for the NetBackup Deduplication Engine:</p> <ul style="list-style-type: none"> <li>■ The media server that hosts the storage server. Although the media server and the storage server share a host, the storage server sends the data through the media server <code>bptm</code> process on that host.</li> <li>■ A load balancing server in the same deduplication node. See <a href="#">“About MSDP load balancing servers”</a> on page 22.</li> <li>■ A deduplication server in a different deduplication node that is the target of optimized duplication. See <a href="#">“About the media servers for MSDP optimized duplication within the same domain”</a> on page 121.</li> </ul> <p>You can specify the server to use for restores. See <a href="#">“Specifying the restore server”</a> on page 248.</p>
Restore directly to the client	<p>The storage server can bypass the media server and move the data directly to the client. NetBackup does not choose a media server for the restore, and the restore does not use the media server <code>bptm</code> process.</p> <p>You must configure NetBackup to bypass a media server and receive the restore data directly from the storage server. See <a href="#">“Configuring MSDP restores directly to a client”</a> on page 247.</p> <p>By default, NetBackup rehydrates the data on the storage server except for client-side deduplication clients. Those clients rehydrate the data. You can configure NetBackup so that the data is rehydrated on the storage server rather than the client. See the <code>RESTORE_DECRYPT_LOCAL</code> parameter in the MSDP <code>pd.conf</code> file. See <a href="#">“MSDP pd.conf file parameters”</a> on page 171. See <a href="#">“Editing the MSDP pd.conf file”</a> on page 170.</p>

## Configuring MSDP restores directly to a client

The NetBackup MSDP storage server can move restore data directly to an MSDP client, bypassing the media server components.

See “[How MSDP restores work](#)” on page 246.

To enable restores directly to a client

- 1 Set the `OLD_VNETD_CALLBACK` option to `YES` on the client. The `OLD_VNETD_CALLBACK` option is stored in the `bp.conf` file on UNIX systems and the registry on Windows systems.

See “[Setting NetBackup configuration options by using the command line](#)” on page 132.

- 2 On the master server, run the following command to configure NetBackup to use client-direct restores for the client:

UNIX: `/usr/openv/netbackup/bin/admincmd/bpclient -client client_name -update -client_direct_restore 2`

Windows: `install_path\NetBackup\bin\admincmd\bpclient -client client_name -update -client_direct_restore 2`

## 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, Volume I*:  
<http://www.symantec.com/docs/DOC5332>

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 7-8** 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 MSDP restores work”](#) on page 246.

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, Volume I*:



<http://www.symantec.com/docs/DOC5332>

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 the *NetBackup Commands Reference Guide*:

<http://www.symantec.com/docs/DOC5332>

# Recovering MSDP

This chapter includes the following topics:

- [About recovering the MSDP catalog](#)
- [Restoring the MSDP catalog from a shadow copy](#)
- [Recovering from an MSDP storage server disk failure](#)
- [Recovering from an MSDP storage server failure](#)
- [Recovering the MSDP storage server after NetBackup catalog recovery](#)

## About recovering the MSDP catalog

The following are the recovery options for the NetBackup MSDP catalog:

**Table 8-1** MSDP catalog backup recovery options

Recovery option	Description
Restore from a shadow copy	<p>If NetBackup detects corruption in the MSDP catalog, the Deduplication Manager restores the catalog automatically from the most recent shadow copy. The automatic restore process also plays a transaction log so that the recovered MSDP catalog is current.</p> <p>Although the shadow copy restore process is automatic, a restore procedure is available if you need to recover from a shadow copy manually.</p> <p>See <a href="#">“About the MSDP shadow catalog”</a> on page 188.</p> <p>See <a href="#">“Restoring the MSDP catalog from a shadow copy”</a> on page 251.</p>

Table 8-1                      MSDP catalog backup recovery options (continued)

Recovery option	Description
Recover from a backup	<p>If you configured an MSDP catalog backup policy and a valid backup exists, you can recover the catalog from a backup. As a general rule, you should only attempt to recover the MSDP catalog from a backup if you have no alternatives. As an example: A hardware problem or a software problem results in the complete loss of the MSDP catalog and the shadow copies.</p> <p>Symantec provides the recovery procedures in the following knowledge base article:</p> <p><a href="http://www.symantec.com/docs/HOWTO95638">http://www.symantec.com/docs/HOWTO95638</a></p> <p>The knowledge base article lets Symantec provide the recovery and the troubleshooting procedures in one location. Symantec can then update the knowledge base article as new information becomes available.</p>

**Caution:** You must determine if your situation is severe enough to recover the catalog. Symantec recommends that you contact your Symantec Support representative before you restore or recover the MSDP catalog. The Support representative can help you determine if you need to recover the catalog or if other solutions are available.

See “[About protecting the MSDP catalog](#)” on page 188.

## Restoring the MSDP catalog from a shadow copy

NetBackup automatically restores the necessary parts of the MSDP catalog if corruption is detected. However, you can restore the MSDP catalog from a shadow copy manually, although in normal circumstances it is not necessary. Symantec recommends that you contact your Symantec Support representative before you restore all or part of the MSDP catalog from a shadow copy.

The procedure that you use depends on the restore scenario, as follows:

- Restore the entire MSDP catalog from a shadow copy

In this scenario, you want to restore the entire catalog from one of the shadow copies.

See “[To restore the entire MSDP catalog from a shadow copy](#)” on page 252.

Restore a specific MSDP database file

The MSDP catalog is composed of multiple small database files. Those files are organized in the file system by the client name and policy name, as follows:

UNIX:

*/database\_path/databases/catalogshadow/2/ClientName/PolicyName*

Windows:

*database\_path\databases\catalogshadow\2\ClientName\PolicyName*

You can restore the database files for a client and a policy combination. The restore of a specific client's and policy's database files is always from the most recent shadow copy.

See ["To restore a specific MSDP database file from a shadow copy"](#) on page 252.

See ["About recovering the MSDP catalog"](#) on page 250.

#### To restore the entire MSDP catalog from a shadow copy

- 1 If any MSDP jobs are active, either cancel them or wait until they complete.
- 2 Disable all policies and storage lifecycle policies that back up to the **Media Server Deduplication Pool**.
- 3 On the MSDP storage server, run the following command, depending on host type:

- UNIX: `/usr/opensv/pdde/pdcr/bin/cacontrol --catalog recover all`
- Windows: `install_path\Veritas\pdde\cacontrol --catalog recover all`

- 4 Enable all policies and storage lifecycle policies that back up to the **Media Server Deduplication Pool**.

- 5 Restart the jobs that were canceled before the recovery.

#### To restore a specific MSDP database file from a shadow copy

- 1 If any MSDP jobs are active for the client and the backup policy combination, either cancel them or wait until they complete.
- 2 Disable the policies and storage lifecycle policies for the client and the backup policy combination that back up to the **Media Server Deduplication Pool**.

- 3 Change to the shadow directory for the client and policy from which you want to recover that database file. That directory contains the database files from which to recover. The following are the pathname formats:

UNIX:

*/database\_path/databases/catalogshadow/2/ClientName/PolicyName*

Windows:

*database\_path\databases\catalogshadow\2\ClientName\PolicyName*

- 4 Run the following command, depending on host type:
  - UNIX: `/usr/opensv/pdde/pdcr/bin/cacontrol --catalog recover 2 "/ClientName/PolicyName"`
  - Windows: `install_path\Veritas\pdde\cacontrol --catalog recover 2 "\ClientName\PolicyName"`
- 5 Enable all policies and storage lifecycle policies that back up to the **Media Server Deduplication Pool**.
- 6 If you canceled jobs before you recovered the database files, restart them.

## Recovering from an MSDP 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 8-2** 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 <i>NetBackup Installation Guide for UNIX and Windows</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a> See "About the MSDP license key" on page 45.
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 an MSDP host configuration file" on page 187.
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 222.
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 222.
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 MSDP storage server configuration" on page 184.
Step 8	Edit the configuration file	See "Editing an MSDP storage server configuration file" on page 184.
Step 9	Configure the storage server	Configure the storage server by uploading the configuration from the file you edited. See "Setting the MSDP storage server configuration" on page 186.
Step 10	Add load balancing servers	If you use load balancing servers in your environment, add them to your configuration. See "Adding an MSDP load balancing server" on page 168.

## Recovering from an MSDP storage server failure

To recover from a permanent failure of the storage server host computer, use the process that is described in this topic.

NetBackup recommends that you consider the following items before you recover:

- The new computer must use the same byte order as the old computer.

---

**Warning:** If the new computer does not use the same byte order as the old computer, 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.

---

- Symantec recommends that the new computer use the same operating system as the old computer.
- Symantec recommends that the new computer use the same version of NetBackup as the old computer.  
If you use a newer version of NetBackup on the new computer, ensure that you perform any data conversions that may be required for the newer release.  
If you want to use an older version of NetBackup on the replacement host, contact your Symantec support representative.

**Table 8-3** Recover from an MSDP 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><b>Warning:</b> Do not delete the images. They are imported back into NetBackup later in this process.</p> <p>If you use the <code>bpexpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>:  <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>
Step 2	Delete the storage units that use the disk pool	<p>See the <i>NetBackup Administrator's Guide, Volume I</i>:  <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>
Step 3	Delete the disk pool	See "Deleting a Media Server Deduplication Pool" on page 234.
Step 4	Delete the deduplication storage server	See "Deleting an MSDP storage server" on page 219.

**Table 8-3** Recover from an MSDP 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 <a href="#">“Deleting an MSDP host configuration file”</a> on page 187.
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 <a href="#">“Deleting credentials from a load balancing server”</a> on page 222.
Step 7	Configure the new host so it meets deduplication requirements	When you configure the new host, consider the following: <ul style="list-style-type: none"> <li>■ You can use the same host name or a different name.</li> <li>■ You can use the same <b>Storage Path</b> or a different <b>Storage Path</b>. If you use a different <b>Storage Path</b>, you must move the deduplication storage to that new location.</li> <li>■ If the <b>Database Path</b> on the original host is different than the <b>Storage Path</b>, you can do one of the following: <ul style="list-style-type: none"> <li>■ You can use the same <b>Database Path</b>.</li> <li>■ You can use a different <b>Database Path</b>. If you do, you must move the deduplication database to the new location.</li> <li>■ You do not have to continue to use a different <b>Database Path</b>. You can move the <code>databases</code> directory into the <b>Storage Path</b> and then specify only the <b>Storage Path</b> when you configure the storage server.</li> </ul> </li> <li>■ You can use the host's default network interface or specify a network interface. If the original host used a specific network interface, you do not have to use the same interface name.</li> </ul> See <a href="#">“About MSDP storage servers”</a> on page 21. See <a href="#">“About MSDP server requirements”</a> on page 22.
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 Windows</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>
Step 10	Reconfigure deduplication	You must use the same credentials for the NetBackup Deduplication Engine.  See <a href="#">“Configuring MSDP server-side deduplication”</a> on page 49.



**Table 8-3** Recover from an MSDP storage server failure (*continued*)

Step	Task	Procedure
Step 11	Import the backup images	See the <i>NetBackup Administrator's Guide, Volume I</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>

# Recovering the MSDP 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 MSDP storage server configuration](#)” on page 186.

Symantec recommends that you save your storage server configuration.

See “[Save the MSDP storage server configuration](#)” on page 37.

Information about recovering the master server is available.

See *NetBackup Troubleshooting Guide*:

<http://www.symantec.com/docs/DOC5332>

# Replacing MSDP hosts

This chapter includes the following topics:

- [Replacing the MSDP storage server host computer](#)

## Replacing the MSDP 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 computer include a lease swap or perhaps the current deduplication storage server computer does not meet your performance requirements.

NetBackup recommends that you consider the following items before you recover:

- The new computer must use the same byte order as the old computer.

---

**Warning:** If the new computer does not use the same byte order as the old computer, 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.

---

- Symantec recommends that the new computer use the same operating system as the old computer.
- Symantec recommends that the new computer use the same version of NetBackup as the old computer.

If you use a newer version of NetBackup on the new computer, ensure that you perform any data conversions that may be required for the newer release.

If you want to use an older version of NetBackup on the replacement host, contact your Symantec support representative.

**Table 9-1** Replacing an MSDP storage server host computer

Step	Task	Procedure
Step 1	Expire the backup images	<p>Expire all backup images that reside on the deduplication disk storage.</p> <p><b>Warning:</b> Do not delete the images. They are imported back into NetBackup later in this process.</p> <p>If you use the <code>bpexpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>:  <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>
Step 2	Delete the storage units that use the disk pool	<p>See the <i>NetBackup Administrator's Guide, Volume I</i>:  <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>
Step 3	Delete the disk pool	See "Deleting a Media Server Deduplication Pool" on page 234.
Step 4	Delete the deduplication storage server	See "Deleting an MSDP storage server" on page 219.
Step 5	Delete the deduplication host configuration file	<p>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.</p> <p>See "Deleting an MSDP host configuration file" on page 187.</p>
Step 6	Delete the credentials on deduplication servers	<p>If you have load balancing servers, delete the NetBackup Deduplication Engine credentials on those media servers.</p> <p>See "Deleting credentials from a load balancing server" on page 222.</p>

**Table 9-1** Replacing an MSDP 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, consider the following:</p> <ul style="list-style-type: none"> <li>■ You can use the same host name or a different name.</li> <li>■ You can use the same <b>Storage Path</b> or a different <b>Storage Path</b>. If you use a different <b>Storage Path</b>, you must move the deduplication storage to that new location.</li> <li>■ If the <b>Database Path</b> on the original host is different than the <b>Storage Path</b>, you can do one of the following: <ul style="list-style-type: none"> <li>■ You can use the same <b>Database Path</b>.</li> <li>■ You can use a different <b>Database Path</b>. If you do, you must move the deduplication database to the new location.</li> <li>■ You do not have to continue to use a different <b>Database Path</b>. You can move the <code>databases</code> directory into the <b>Storage Path</b> and then specify only the <b>Storage Path</b> when you configure the storage server.</li> </ul> </li> <li>■ You can use the host's default network interface or specify a network interface.</li> </ul> <p>If the original host used a specific network interface, you do not have to use the same interface name.</p> <p>See “<a href="#">About MSDP storage servers</a>” on page 21.</p> <p>See “<a href="#">About MSDP server requirements</a>” on page 22.</p>
Step 8	Connect the storage to the host	<p>Use the storage path that you configured for this replacement host.</p> <p>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 Windows</i>:</p> <p><a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>
Step 10	Reconfigure deduplication	See “ <a href="#">Configuring MSDP server-side deduplication</a> ” on page 49.
Step 11	Import the backup images	<p>See the <i>NetBackup Administrator's Guide, Volume I</i>:</p> <p><a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>

# Uninstalling MSDP

This chapter includes the following topics:

- [About uninstalling MSDP](#)
- [Deactivating MSDP](#)

## About uninstalling MSDP

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 MSDP storage server name or storage path”](#) on page 216.
- Deactivate deduplication and remove the configuration files and the storage files.  
See [“Deactivating MSDP”](#) on page 261.

## Deactivating MSDP

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 10-1** Remove MSDP

Step	Task	Procedure
Step 1	Remove client deduplication	Remove the clients that deduplicate their own data from the client deduplication list.  See <a href="#">"Disabling MSDP client-side deduplication for a client"</a> on page 107.
Step 2	Delete the storage units that use the disk pool	See the <i>NetBackup Administrator's Guide, Volume I</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>
Step 3	Delete the disk pool	See <a href="#">"Deleting a Media Server Deduplication Pool"</a> on page 234.
Step 4	Delete the deduplication storage server	See <a href="#">"Deleting an MSDP storage server"</a> on page 219.  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 <a href="#">"Deleting the MSDP storage server configuration"</a> on page 220.
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 <a href="#">"Deleting an MSDP host configuration file"</a> on page 187.
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.)  <b>Warning:</b> 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:

- MSDP server components
- Media server deduplication backup process
- MSDP client components
- MSDP client-side deduplication backup process

## MSDP server components

Figure 11-1 is a diagram of the storage server components.

Figure 11-1 MSDP server components

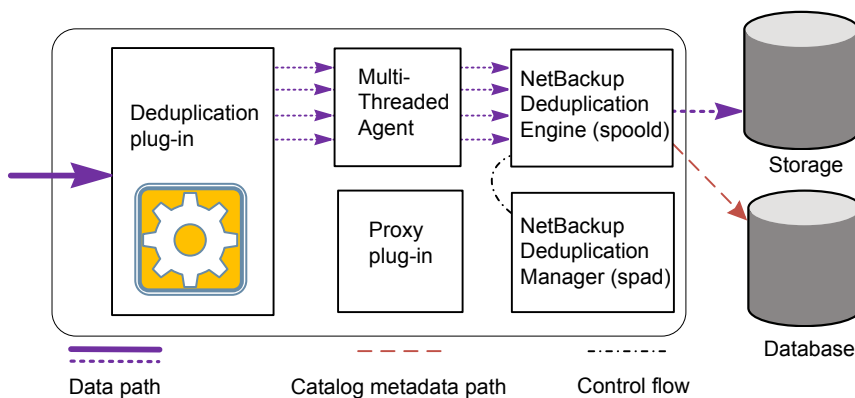


Table 11-1 describes the MSDP server components.

**Table 11-1** NetBackup MSDP server components

Component	Description
Deduplication plug-in	<p>The deduplication plug-in does the following:</p> <ul style="list-style-type: none"> <li>■ Separates the file's metadata from the file's content.</li> <li>■ Deduplicates the content (separates files into segments).</li> <li>■ If required, compresses the data for backups and decompresses the backups for restores.</li> <li>■ If required, encrypts the data for backups and decrypts the backups for restores.</li> <li>■ If required, compresses the data for duplication and replication transfer.</li> <li>■ If required, encrypts the data for duplication and replication transfer.</li> </ul> <p>The plug-in runs on the deduplication storage server and on load balancing servers.</p>
Multi-Threaded Agent	<p>The NetBackup Deduplication Multi-Threaded Agent uses multiple threads for asynchronous network I/O and CPU core calculations. The agent runs on the storage server, load balancing servers, and clients that deduplicate their own data.</p> <p>See <a href="#">"About the MSDP Deduplication Multi-Threaded Agent"</a> on page 52.</p>
NetBackup Deduplication Engine	<p>The NetBackup Deduplication Engine is one of the storage server core components. It provides many of the deduplication functions, which are described in <a href="#">Table 11-2</a>.</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>
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>
Proxy plug-in	<p>The proxy plug-in manages control communication with the clients that back up their own data. It communicates with the OpenStorage proxy server (<code>nbostrpxy</code>) on the client.</p>
Reference database	<p>The reference database stores the references that point to every data segment of which a file is composed. Unique fingerprints identify data segments. The reference database is partitioned into multiple small reference database files to improve scalability and performance.</p> <p>The reference database is separate from the NetBackup catalog. The NetBackup catalog maintains the usual NetBackup backup image information.</p>

[Table 11-2](#) describes the components and functions within the NetBackup Deduplication Engine.



**Table 11-2** NetBackup Deduplication Engine components and functions

Component	Description
Connection and Task Manager	<p>The Connection and Task Manager manages all of the connections from the load balancing servers and the clients that deduplicate their own data. The Connection and Task Manager is a set of functions and threads that does the following:</p> <ul style="list-style-type: none"><li>■ Provides a thread pool to serve all clients.</li><li>■ Maintains a task for each client connection.</li><li>■ Manages the mode of the Deduplication Engine based on the operation. Operations are backups, restores, queue processing, and so on.</li></ul>
Data integrity checking	<p>The NetBackup Deduplication Engine checks the integrity of the data and resolves integrity problems.</p> <p>See <a href="#">“About MSDP data integrity checking”</a> on page 236.</p>
Data Store Manager	<p>The Data Store Manager manages all of the data container files. The datastore Manager is a set of functions and threads that provides the following:</p> <ul style="list-style-type: none"><li>■ A transaction mechanism to back up data into the datastore.</li><li>■ A mechanism to read data from the datastore.</li><li>■ A transaction mechanism to reclaim space in the datastore (that is, compact containers and remove containers). Container IDs are unique. The Data Store Manager increments the container number with each new container created. The data in a container is never overwritten, and a container ID is never reused.</li></ul> <p>See <a href="#">“About MSDP container files”</a> on page 205.</p>
Index Cache Manager	<p>The Index Cache Manager manages the fingerprint cache. The cache improves fingerprint lookup speed.</p> <p>See <a href="#">“About the MSDP fingerprint cache”</a> on page 60.</p>
Queue processing	<p>The NetBackup Deduplication Engine processes the transaction queue.</p> <p>See <a href="#">“About MSDP queue processing”</a> on page 235.</p>
Reference Database Engine	<p>The Reference Database Engine stores the references that point to the data segments, such as read-from or write-to references. It manipulates a single database file at a time.</p>

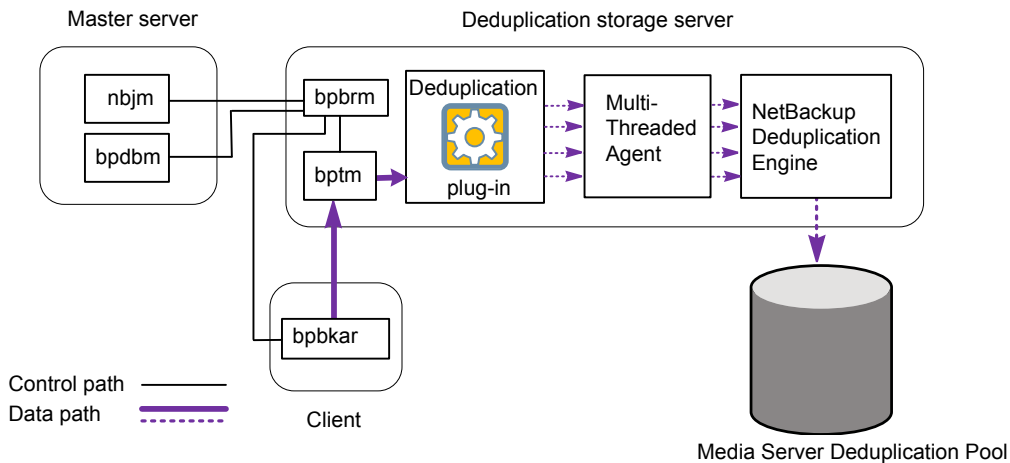
**Table 11-2** NetBackup Deduplication Engine components and functions  
(continued)

Component	Description
Reference Database Manager	The Reference Database Manager manages all of the container references. It provides a transaction mechanism to manipulate a single database file.

## Media server deduplication backup process

The [Figure 11-2](#) diagram shows the backup process when a media server deduplicates the backups. The destination is a **Media Server Deduplication Pool**. A description follows.

**Figure 11-2** 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 (`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`) on the client 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 IDs of the container files from the NetBackup Deduplication Engine. Those container files contain the fingerprints from the last full backup for the client. 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 separates the files in the backup image into segments.
  - The deduplication plug-in buffers the segments and then sends batches of them to the Deduplication Multi-Threaded Agent. Multiple threads and shared memory are used for the data transfer.
  - The NetBackup Deduplication Multi-Threaded Agent processes the data segments in parallel using multiple threads to improve throughput performance. The agent then sends only the unique data segments to the NetBackup Deduplication Engine.
- If the host is a load balancing server, the Deduplication Engine is on a different host, the storage server.
- The NetBackup Deduplication Engine writes the data to the **Media Server Deduplication Pool**.

The first backup may have a 0% deduplication rate, although a 0% rate is unlikely. Zero percent means that all file segments in the backup data are unique.


The [Figure 11-3](#) diagram shows the backup process when a media server deduplicates the backups. The destination is a **PureDisk Deduplication Pool**. A description follows.

---

**Note:** NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

---

**Figure 11-3** 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 IDs of the container files from the NetBackup Deduplication Engine. Those container files contain the fingerprints from the last full backup for the client. 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 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**.

# MSDP client components

Table 11-3 describes the client deduplication components.


**Table 11-3** Client MSDP components

Component	Description
Deduplication plug-in	The deduplication plug-in does the following: <ul style="list-style-type: none"><li>■ Separates the file's metadata from the file's content.</li><li>■ Deduplicates the content (separates files into segments).</li><li>■ If required, compresses the data for backups and decompresses the backups for restores.</li><li>■ If required, encrypts the data for backups and decrypts the backups for restores.</li></ul>
Multi-Threaded Agent	The NetBackup Deduplication Multi-Threaded Agent uses multiple threads for asynchronous network I/O and CPU core calculations. The agent runs on the storage server, load balancing servers, and clients that deduplication their own data.  See <a href="#">“About the MSDP Deduplication Multi-Threaded Agent”</a> on page 52.
Proxy server	The OpenStorage proxy server ( <code>nboostpxy</code> ) manages control communication with the proxy plug-in on the storage server.

## MSDP client–side deduplication backup process

The [Figure 11-4](#) diagram shows the backup process of a client that deduplicates its own data. The destination is a media server deduplication pool. A description follows.

**Figure 11-4** MSDP client backup to a deduplication pool



The following list describes the backup process for an MSDP 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 on the client tries to retrieve a list of fingerprints, in the following order:
  - From a client and a policy that is configured in the client's `pd.conf` file. The `FP_CACHE_CLIENT_POLICY` entry defines the client and policy to use for the fingerprint cache. The entry must be valid (that is, not expired).  
See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 62.
  - From the previous backup for the client and policy.
  - From the special seeding directory on the storage server.  
See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 62.

The list of fingerprints is used as a cache so the plug-in does not have to request each fingerprint from the engine.

If no fingerprints are loaded into the cache, the deduplication rate may be very low for the backup.

- The deduplication plug-in separates the files in the backup image into segments.
- The deduplication plug-in buffers the segments and then sends batches of them to the Deduplication Multi-Threaded Agent. Multiple threads and shared memory are used for the data transfer.
- The NetBackup Deduplication Multi-Threaded Agent processes the data segments in parallel using multiple threads to improve throughput performance. The agent then sends only the unique data segments to the NetBackupDeduplication Engine.
- The NetBackup Deduplication Engine writes the data to the **Media Server Deduplication Pool**.  
The first backup may have a 0% deduplication rate, although a 0% deduplication rate is unlikely. Zero percent means that all file segments in the backup data are unique.


The [Figure 11-5](#) diagram shows the backup process of a client that deduplicates its own data. The destination is a PureDisk storage pool. A description follows.

---

**Note:** NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

---

**Figure 11-5** MSDP deduplication client backup to a PureDisk storage pool



The following list describes the backup process for an MSDP client to a PureDisk storage 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 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 (bpbm). The Backup/Restore Manager sends the file information to the bpbm 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 of the client. 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 fingerprint calculations.
- The deduplication plug-in sends only unique data segments to the PureDisk storage pool.

# Troubleshooting

This chapter includes the following topics:

- [About unified logging](#)
- [About legacy logging](#)
- [About MSDP log files](#)
- [Troubleshooting MSDP installation issues](#)
- [Troubleshooting MSDP configuration issues](#)
- [Troubleshooting MSDP operational issues](#)
- [Viewing MSDP disk errors and events](#)
- [MSDP event codes and messages](#)

## About unified logging

Unified logging and legacy logging are the two forms of debug logging used in NetBackup. Unified logging creates log file names and messages in a standardized format. All NetBackup processes use either unified logging or legacy logging.

Unlike the files that are written in legacy logging, unified logging files cannot be viewed with a text editor. The unified logging files are in binary format, and some of the information is contained in an associated resource file.

See [“About legacy logging”](#) on page 277.

Server processes and client processes use unified logging.

Unlike legacy logging, unified logging does not require that you create logging subdirectories. Log files for originator IDs are written to a subdirectory with the name specified in the log configuration file. All unified logs are written to subdirectories in the following directory:

Windows `install_path\NetBackup\logs`

UNIX `/usr/opensv/logs`

You can access logging controls in the **NetBackup Administration Console**. In the left pane, expand **NetBackup Management > Host Properties > Master Servers** or **Media Servers**. Double-click the server you want to change. In the left pane of the dialog box, click **Logging**.

You can also manage unified logging by using the following commands:

<code>vxlogcfg</code>	Modifies the unified logging configuration settings. for more information about the <code>vxlogcfg</code> command.
<code>vxlogmgr</code>	Manages the log files that the products that support unified logging generate. for more information about the <code>vxlogmgr</code> command.
<code>vxlogview</code>	Displays the logs that unified logging generates. See <a href="#">“Examples of using vxlogview to view unified logs”</a> on page 276. for more information about the <code>vxlogview</code> command.

See the *NetBackup Commands Reference Guide* for a complete description about these commands. The guide is available through the following URL:

<http://www.symantec.com/docs/DOC5332>

These commands are located in the following directory:

Windows `install_path\NetBackup\bin`

UNIX `/usr/opensv/netbackup/bin`

## About using the vxlogview command to view unified logs

Use the `vxlogview` command to view the logs that unified logging creates. These logs are stored in the following directory.

UNIX `/usr/opensv/logs`

Windows `install_path\NetBackup\logs`

Unlike the files that are written in legacy logging, unified logging files cannot be easily viewed with a text editor. The unified logging files are in binary format, and

some of the information is contained in an associated resource file. Only the `vxlogview` command can assemble and display the log information correctly.

You can use `vxlogview` to view NetBackup log files as well as PBX log files.

To view PBX logs using the `vxlogview` command, do the following:

- Ensure that you are an authorized user. For UNIX and Linux, you must have root privileges. For Windows, you must have administrator privileges.
- To specify the PBX product ID, enter `-p 50936` as a parameter on the `vxlogview` command line.

`vxlogview` searches all the files, which can be a slow process. Refer to the following topic for an example of how to display results faster by restricting the search to the files of a specific process.

## Examples of using vxlogview to view unified logs

The following examples demonstrate how to use the `vxlogview` command to view unified logs.

**Table 12-1** Example uses of the `vxlogview` command

Item	Example
Display all the attributes of the log messages	<code>vxlogview -p 51216 -d all</code>
Display specific attributes of the log messages	Display the log messages for NetBackup (51216) that show only the date, time, message type, and message text:  <code>vxlogview --prodid 51216 --display D,T,m,x</code>
Display the latest log messages	Display the log messages for originator 116 ( <code>nbpem</code> ) that were issued during the last 20 minutes. Note that you can specify <code>-o nbpem</code> instead of <code>-o 116</code> :  <code># vxlogview -o 116 -t 00:20:00</code>
Display the log messages from a specific time period	Display the log messages for <code>nbpem</code> that were issued during the specified time period:  <code># vxlogview -o nbpem -b "05/03/05 06:51:48 AM" -e "05/03/05 06:52:48 AM"</code>

**Table 12-1** Example uses of the vxlogview command (*continued*)

Item	Example
Display results faster	<p>You can use the <code>-i</code> option to specify an originator for a process:</p> <pre># vxlogview -i nbpem</pre> <p>The <code>vxlogview -i</code> option searches only the log files that the specified process (<code>nbpem</code>) creates. By limiting the log files that it has to search, <code>vxlogview</code> returns a result faster. By comparison, the <code>vxlogview -o</code> option searches all unified log files for the messages that the specified process has logged.</p> <p><b>Note:</b> If you use the <code>-i</code> option with a process that is not a service, <code>vxlogview</code> returns the message "No log files found." A process that is not a service has no originator ID in the file name. In this case, use the <code>-o</code> option instead of the <code>-i</code> option.</p> <p>The <code>-i</code> option displays entries for all OIDs that are part of that process including libraries (137, 156, 309, etc.).</p>
Search for a job ID	<p>You can search the logs for a particular job ID:</p> <pre># vxlogview -i nbpem   grep "jobid=job_ID"</pre> <p>The <code>jobid=</code> search key should contain no spaces and must be lowercase.</p> <p>When searching for a job ID, you can use any <code>vxlogview</code> command option. This example uses the <code>-i</code> option with the name of the process (<code>nbpem</code>). The command returns only the log entries that contain the job ID. It misses related entries for the job that do not explicitly contain the <code>jobid=job_ID</code>.</p>

See the *NetBackup Commands Reference Guide* for a complete description of the `vxlogview` command. The guide is available through the following URL:

<http://www.symantec.com/docs/DOC5332>

## About legacy logging

Legacy logging and unified logging are the two forms of debug logging used in NetBackup. In legacy debug logging, each process creates logs of debug activity in its own logging directory. All NetBackup processes use either unified logging or legacy logging.

See “[About unified logging](#)” on page 274.

To enable legacy debug logging on NetBackup servers, you must first create the appropriate directories for each process.

Windows      *install\_path\NetBackup\logs*  
                  *install\_path\Volmgr\debug*

UNIX            */usr/opensv/netbackup/logs*  
                  */usr/opensv/volmgr/debug*

After the directories are created, NetBackup creates log files in the directory that are associated with each process. A debug log file is created when the process begins.

To enable debug logging for the NetBackup Status Collection Daemon (*vmacd*), create the following directory before you start *nbemm*.

Windows      *install\_path\Volmgr\debug\reqlib\*  
 UNIX            */usr/opensv/volmgr/debug/reqlib*

As an alternative, you can stop and restart *nbemm* after creating the directory:

Tables are available that list the log directories that you must create.

You can use the following batch file to create all the debug log directories at once.

Windows: *install\_path\NetBackup\Logs\mklogdir.bat*

UNIX: *usr/opensv/netbackup/logs/mklogdir*

Media servers have only the *bpbm*, *bpcd*, *bpdm*, and *bptm* debug logs.

## Creating NetBackup log file directories

Before you configure your NetBackup feature, create the directories into which NetBackup commands write log files. Create the directories on the master server and on each media server that you use for your feature. The log files reside in the following directories:

- UNIX: */usr/opensv/netbackup/logs/*
- Windows: *install\_path\NetBackup\logs\*

More information about NetBackup logging is available in the *NetBackup Troubleshooting Guide*, available through the following URL:

<http://www.symantec.com/docs/DOC5332>

### To create log directories for NetBackup commands

- ◆ Depending on the operating system, run one of the following scripts:

UNIX: `/usr/opensv/netbackup/logs/mklogdir`

Windows: `install_path\NetBackup\logs\mklogdir.bat`

### To create the `tpconfig` command log directory

- ◆ Depending on the operating system, create the `debug` directory and the `tpcommand` directory (by default, the `debug` directory and the `tpcommand` directory do not exist). The pathnames of the directories are as follows:

UNIX: `/usr/opensv/volmgr/debug/tpcommand`

Windows: `install_path\Veritas\Volmgr\debug\tpcommand`

## About MSDP log files

The NetBackup deduplication components write information to various log files. Some NetBackup commands or processes write messages to their own log files. Other processes use the 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.

See “[About legacy logging](#)” on page 277.

See “[About unified logging](#)” on page 274.

In VxUL logs, the messages that begin with an `sts` prefix relate to the interaction with the deduplication 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 Logging Reference Guide*. The guide is available through the following URL:

<http://www.symantec.com/docs/DOC5332>

Most interaction occurs on the NetBackup media servers. Therefore, the log files on the media servers that you use for disk operations are of most interest.

---

**Warning:** The higher the log level, the greater the affect on NetBackup performance. Use a log level of 5 (the highest) only when directed to do so by a Symantec representative. A log level of 5 is for troubleshooting only.

Specify the NetBackup log levels in the **Logging** host properties on the NetBackup master server. The log levels for some processes specific to certain options are set in configuration files as described in [Table 12-2](#).

---

Table 12-2 describes the log files for each component.

**Table 12-2** Logs for NetBackup MSDP activity

Component	VxUL OID	Description
Backups and restores	117	The nbjm Job Manager.
Backups and restores	N/A	<p>Messages appear in the log files for the following processes:</p> <ul style="list-style-type: none"> <li>■ The bpbrm backup and restore manager. The following is the path to the log files:            UNIX: <code>/usr/opensv/netbackup/logs/bpbrm</code>            Windows: <code>install_path\Veritas\NetBackup\logs\bpbrm</code> </li> <li>■ The bpdgm database manager. The following is the path to the log files:            UNIX: <code>/usr/opensv/netbackup/logs/bpdgm</code>            Windows: <code>install_path\Veritas\NetBackup\logs\bpdgm</code> </li> <li>■ The bptm tape manager for I/O operations. The following is the path to the log files:            UNIX: <code>/usr/opensv/netbackup/logs/bptm</code>            Windows: <code>install_path\Veritas\NetBackup\logs\bptm</code> </li> </ul>
Catalog shadow copies	N/A	<p>The MSDP catalog shadow copy process writes messages to the following log files and directories:</p> <p>UNIX:</p> <pre> /storage_path/log/spad/spad.log /storage_path/log/spad/sched_CatalogBackup.log /storage_path/log/spad/client_name/ </pre> <p>Windows:</p> <pre> storage_path\log\spad\spad.log storage_path\log\spad\sched_CatalogBackup.log storage_path\log\spad\client_name\ </pre>



**Table 12-2** Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
Client deduplication proxy plug-in	N/A	<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"> <li>■ For <code>bptm</code>:            UNIX: <code>/usr/openv/netbackup/logs/bptm</code>            Windows: <code>install_path\Veritas\NetBackup\logs\bptm</code> </li> <li>■ For <code>bpstsinfo</code>:            Windows: <code>/usr/openv/netbackup/logs/admin</code>            UNIX: <code>/usr/openv/netbackup/logs/bpstsinfo</code>            Windows: <code>install_path\Veritas\NetBackup\logs\admin</code>            Windows: <code>install_path\Veritas\NetBackup\logs\stsinfo</code> </li> <li>■ For <code>bpbrm</code>:            UNIX: <code>/usr/openv/netbackup/logs/bpbrm</code>            Windows: <code>install_path\Veritas\NetBackup\logs\bpbrm</code> </li> </ul>
Client deduplication proxy server	N/A	<p>The deduplication proxy server <code>nbostrpxy</code> on the client writes messages to files in the following directory, as follows:</p> <p>UNIX: <code>/usr/openv/netbackup/logs/nbostrpxy</code></p> <p>Windows: <code>install_path\Veritas\NetBackup\logs\nbostrpxy</code>.</p>
Deduplication configuration script	N/A	<p>The following is the path name of the log file for the deduplication configuration script:</p> <ul style="list-style-type: none"> <li>■ UNIX: <code>storage_path/log/pdde-config.log</code></li> <li>■ Windows: <code>storage_path\log\pdde-config.log</code></li> </ul> <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 plug-in	N/A	<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 <a href="#">“About the MSDP pd.conf configuration file”</a> on page 169.</p> <p>See <a href="#">“Editing the MSDP pd.conf file”</a> on page 170.</p>
Device configuration and monitoring	111	The <code>nbeemm</code> process.

**Table 12-2** Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
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.
<code>drcontrol</code> utility	N/A	<p>You must run the <code>drcontrol</code> utility on the MSDP storage server host. The command requires administrator privileges.</p> <p>The utility creates a log file and displays its pathname in the command output. The utility writes log files to the following directory, depending on the operating system:</p> <p>UNIX:</p> <pre>[storage_path]/log/drcontrol/policy_admin</pre> <pre>/storage_path/log/drcontrol/dedupe_catalog_DR</pre> <p>Windows:</p> <pre>storage_path\log\drcontrol\policy_admin</pre> <pre>storage_path\log\drcontrol\dedupe_catalog_DR</pre> <p>See <a href="#">“About protecting the MSDP catalog”</a> on page 188.</p> <p>See <a href="#">“About recovering the MSDP catalog”</a> on page 250.</p>
Installation	N/A	<p>The NetBackup installation process writes information about the installation of the deduplication components to a log file in the following directory:</p> <ul style="list-style-type: none"> <li>■ UNIX: <code>/var/log/puredisk</code></li> <li>■ Windows: <code>%ALLUSERSPROFILE%\Symantec\NetBackup\InstallLogs</code></li> </ul>

**Table 12-2** Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
NetBackup Deduplication Engine	N/A	<p>The NetBackup Deduplication Engine writes several log files, as follows:</p> <ul style="list-style-type: none"> <li>Log files in the <i>storage_path/log/spoold</i> directory, as follows: <ul style="list-style-type: none"> <li>The <i>spoold.log</i> file is the main log file</li> <li>The <i>stored.log</i> file is for queue processing.</li> <li>A log file for each connection to the engine is stored in a directory in the storage path <i>spoold</i> directory. The following describes the pathname to a log file for a connection: <i>hostname/application/TaskName/MMDDYY.log</i> For example, the following is an example of a <i>crcontrol</i> connection log pathname on a Linux system: <i>/storage_path/log/spoold/server./crcontrol/Control/010112.log</i> Usually, the only reason to examine these connection log files is if a Symantec support representative asks you to.</li> </ul> </li> <li>A VxUL log file for the events and errors that NetBackup receives from polling. The originator ID for the deduplication engine is 364.</li> </ul>
NetBackup Deduplication Engine	364	<p>The NetBackup Deduplication Engine that runs on the deduplication storage server.</p>
NetBackup Deduplication Manager	N/A	<p>The log files are in the <i>/storage_path/log/spad</i> directory, as follows:</p> <ul style="list-style-type: none"> <li><i>spad.log</i></li> <li><i>sched_QueueProcess.log</i></li> <li><i>SchedClass.log</i></li> <li>A log file for each connection to the manager is stored in a directory in the storage path <i>spad</i> directory. The following describes the pathname to a log file for a connection: <i>hostname/application/TaskName/MMDDYY.log</i> For example, the following is an example of a <i>bpstsinfo</i> connection log pathname on a Linux system: <i>/storage_path/log/spoold/server./bpstsinfo/spad/010112.log</i> Usually, the only reason to examine these connection log files is if a Symantec support representative asks you to.</li> </ul> <p>You can set the log level and retention period in the <b>Change Storage Server</b> dialog box <b>Properties</b> tab.</p> <p>See <a href="#">“Changing MSDP storage server properties”</a> on page 214.</p>

**Table 12-2** Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
Optimized duplication and replication	N/A	<p>For optimized duplication and Auto Image Replication, The following are the log files that provide information:</p> <ul style="list-style-type: none"> <li>■ The NetBackup <code>bptm</code> tape manager for I/O operations. The following is the path to the log files:            UNIX: <code>/usr/openv/netbackup/logs/bptm</code>            Windows: <code>install_path\Veritas\NetBackup\logs\bptm</code></li> <li>■ The following is the path name of MSDP replication log file:  <code>/storage_path/log/spad/replication.log</code></li> </ul>
Resilient network connections	387	<p>The Remote Network Transport Service (<code>nbrntd</code>) manages resilient network connection sockets. It runs on the master server, on media servers, and on clients. Use the VxUL originator ID 387 to view information about the socket connections that NetBackup uses.</p> <p><b>Note:</b> If multiple backup streams run concurrently, the Remote Network Transport Service writes a large amount of information to the log files. In such a scenario, Symantec recommends that you set the logging level for OID 387 to 2 or less. To configure unified logs, see the following guide:</p> <p>The <i>NetBackup Troubleshooting Guide</i>:  <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></p>
Resilient network connections	N/A	<p>The deduplication plug-in logs information about keeping the connection alive. For more information about the deduplication plug-in log file, see “Deduplication plug-in” in this table.</p> <p>The <code>pd.conf</code> file <code>FILE_KEEP_ALIVE_INTERVAL</code> parameter controls the connection keep alive interval.</p> <p>See “<a href="#">About the MSDP pd.conf configuration file</a>” on page 169.</p> <p>See “<a href="#">Editing the MSDP pd.conf file</a>” on page 170.</p>

## Troubleshooting MSDP installation issues

The following sections may help you troubleshoot installation issues.

See “[MSDP installation on SUSE Linux fails](#)” on page 284.

### MSDP 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 MSDP configuration issues

The following sections may help you troubleshoot configuration issues.

See [“About MSDP log files”](#) on page 279.

See [“MSDP storage server configuration fails”](#) on page 285.

See [“MSDP database system error \(220\)”](#) on page 286.

See [“MSDP server not found error”](#) on page 286.

See [“License information failure during MSDP configuration”](#) on page 287.

See [“The disk pool wizard does not display an MSDP volume”](#) on page 288.

### MSDP 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 an MSDP host configuration file”](#) on page 187.

## MSDP 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 MSDP log files”](#) on page 279..

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.

## MSDP 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** RDSM has encountered an issue with STS where the server was not found: `getStorageServerInfo`

```
Failed to create storage server ssname, invalid
command parameter
```

#### Diagnosis

#### Possible root causes:

- 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 `nbdevconfig` 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 MSDP 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 an MSDP 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 MSDP operational issues

The following sections may help you troubleshoot operational issues.

See [“Verify that the MSDP server has sufficient memory”](#) on page 289.

See [“MSDP backup or duplication job fails”](#) on page 289.

See [“MSDP client deduplication fails”](#) on page 291.

See [“MSDP volume state changes to DOWN when volume is unmounted”](#) on page 291.

See [“MSDP errors, delayed response, hangs”](#) on page 292.

See [“Cannot delete an MSDP disk pool”](#) on page 293.

See [“MSDP media open error \(83\)”](#) on page 294.

See [“MSDP media write error \(84\)”](#) on page 296.

See [“MSDP no images successfully processed \(191\)”](#) on page 299.

See [“MSDP storage full conditions”](#) on page 299.

See [“Troubleshooting MSDP catalog backup”](#) on page 300.



## Verify that the MSDP 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 MSDP server requirements”](#) on page 22.

If the NetBackup deduplication processes do not start on Red Hat Linux, configure shared memory to be at least 128 MB (`SHMMAX=128MB`).

## MSDP backup or duplication job fails

The following subsections describe some potential failures for backup or deduplication jobs and how to resolve them.

- [Disk volume is down](#)
- [Storage server is down or unavailable](#)
- [Backup job: System error occurred \(174\)](#)
- [Failure to open storage path or to prepare CRQP transaction](#)

### Disk volume is down

A message similar to the following appears in the job details:

```
Error 800: Disk Volume is Down
```

Examine the disk error logs to determine why the volume was marked DOWN.

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.

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.

See [“Viewing MSDP disk errors and events”](#) on page 301.

### Storage server is down or unavailable

Windows servers only.

A message similar to the following appears in the job details:

```
Error nbjm(pid=6384) NBU status: 2106, EMM status: Storage Server is down or unavailable Disk storage server is down(2106)
```

The NetBackup Deduplication Manager (`spad.exe`) and the NetBackup Deduplication Engine (`spoold.exe`) 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.

To resolve the issue, specify the following shared memory value in the configuration file:

```
SharedMemoryEnabled=1
```

Then, restart both components. Do not change the values of the other two shared memory parameters.

The `SharedMemoryEnabled` parameter is stored in the following file:

```
storage_path\etc\puredisk\agent.cfg
```

## Backup job: System error occurred (174)

A message similar to the following appears in the job details:

```
media manager - system error occurred (174)
```

If the job details also include errors similar to the following, it indicates that an image clean-up job failed:

```
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
```

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**.

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.

See [“About MSDP queue processing”](#) on page 235.

See [“About MSDP log files”](#) on page 279.

## Failure to open storage path or to prepare CRQP transaction

Error messages similar to the following appear in one of the NetBackup Deduplication Engine (`spoold`) log files.

```
RefDBEngine::write_prepare fail to open
/storage_path/databases/refdb/prepare/64.ref.prepare
```

```
RefDBManager::write_prepare fail to prepare CRQP transaction for
refdb 64
```

See [“About MSDP log files”](#) on page 279.

This error occurs if the `/storage_path/databases/refdb/prepare` directory is deleted.

To fix this problem, do one of the following:

- Create the missing directory manually.
- Restart the NetBackup Deduplication Engine (`spoold`). First ensure that no backups are running on the storage unit on that media server.

---

**Note:** `RefDBEngine` and `refdb` do not refer to nor are they related to the open source RefDB reference database and bibliography tool.

---

## MSDP 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 33.

## MSDP 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/opensv/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       : @aaaaab
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

- 1 Mount the file system

After a brief period of time, the volume state changes to UP. No further action is required.

- 2 If the volume state does not change, change it manually.

See [“Changing the MSDP disk volume state”](#) on page 232.

## MSDP errors, delayed response, hangs

Insufficient memory or inadequate host capabilities may cause multiple errors, delayed response, and hangs.

See [“About MSDP server requirements”](#) on page 22.

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 an MSDP 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 MSDP disk](#)
- [Incomplete SLP duplication jobs](#)

### Expired fragments remain on MSDP 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:** `/usr/opensv/netbackup/bin/admincmd/nbstlutil stlilist -incomplete`

**Windows:** `install_path\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:** `/usr/opensv/netbackup/bin/admincmd/nbstlutil cancel -backupid xxxxx`

**Windows:** `install_path\NetBackup\bin\admincmd\nbstlutil cancel -backupid xxxxx`

## MSDP media open error (83)

The `media open error (83)` message is a generic error for the duplication. The error appears in the **NetBackup Administration Console Activity Monitor**.

Often, the NetBackup Deduplication Engine (`spoold`) or the NetBackup Deduplication Manager (`spad`) were too busy to respond to the deduplication process in a timely manner. External factors may cause the Deduplication Engine or the Deduplication Manager to be unresponsive. Were they temporarily busy (such as queue processing in progress)? Do too many jobs run concurrently?

See [“About MSDP performance”](#) on page 31.

Usually but not always the NetBackup `bpdm` log provides additional information about status 83.

The following subsections describe use cases that generated an error 83.

### SQL Server client-side backups fail

Client-side backups of a SQL Server database may fail in the following circumstances:

- The **Both IPv4 and IPv6** option is enabled for the master server, the media server that hosts the NetBackup Deduplication Engine, and the client. The **Both IPv4 and IPv6** option is configured in the **Network Settings** host properties.

- The IPv6 network is configured as a preferred network for the master server, the media server that hosts the NetBackup Deduplication Engine, and the client. The preferred network **Match (Above network will be preferred for communication)** property also is enabled. Preferred networks are configured in the **Preferred Networks** host properties.
- The IPv6 network is chosen for the backup.

Examine the `bpbrm` log file for an error similar to the following:

```
probe_ost_plugin: sts_get_server_prop_byname failed: error 2060057
```

If the error message appears, the NetBackup host name cache may not contain the correct host name mapping information. The cache may be out of sync if DNS changes in your network environment were not fully propagated throughout your environment. It takes some amount of time for DNS changes to propagate throughout a network environment.

To resolve the problem, do the following on the NetBackup master server and on the MSDP storage server:

1. Stop the NetBackup services.
2. Run the following command:

UNIX: `/usr/opensv/netbackup/bin/bpclntcmd -clearhostcache`

Windows: `install_path\NetBackup\bin\bpclntcmd.exe -clearhostcache`

3. Start the NetBackup services.

For more information about client deduplication logging, see the description of “Client deduplication proxy plug-in” in the “MSDP log files” topic.

See [“About MSDP log files”](#) on page 279.

## Restore or duplication fails

The `media open error (83)` message appears in the **NetBackup Administration Console Activity Monitor**.

[Table 12-3](#) describes other messages that may appear.

**Table 12-3** Case sensitivity error messages

Operation	Activity Monitor job details	Status in <code>bpdm</code> and <code>bptm</code> log files
Restore	Image open failed: error 2060018: file not found	<code>sts_open_image</code> failed: error 2060018

**Table 12-3** Case sensitivity error messages (*continued*)

Operation	Activity Monitor job details	Status in <code>bpdm</code> and <code>bptm</code> log files
Duplication (MSDP source)	Image open failed: error 2060018: file not found	<code>sts_open_image</code> failed: error 2060018
Replication (MSDP source)	get image properties failed: error 2060013: no more entries	<code>rpl_add_image_set:</code> <code>rpl_get_image_info()</code> failed, error 2060013
Duplication (NetBackup 5000 series appliance source)	Image open failed: error 2060018: file not found	<code>sts_open_image</code> failed: error 2060018
Replication (NetBackup 5000 series appliance source)	get image properties failed: error 2060013: no more entries	<code>rpl_add_image_set:</code> <code>rpl_get_image_info()</code> failed, error 2060013

The messages may indicate a client name case sensitivity issue in your MSDP environment. For a description of the problem and the procedures to resolve it, see the following Symantec tech note:

<http://www.symantec.com/docs/TECH207194>

## MSDP media write error (84)

[Table 12-4](#) describes solutions to the media write errors that may occur during **Media Server Deduplication Pool** backups, duplication, or replication.

Also see the following subsections for descriptions of more complicated solutions:

- [Host name resolution problems](#)

**Table 12-4** Media write error causes

The NetBackup Deduplication Engine ( <code>spoold</code> ) was too busy to respond.	Examine the Disk Logs report for errors that include the name PureDisk. Examine the disk monitoring services log files for details from the deduplication plug-in.  See " <a href="#">Viewing MSDP disk reports</a> " on page 207.
---	--



**Table 12-4** Media write error causes (*continued*)

Data removal is running.	Data cannot be backed up at the same time as it is removed.  See <a href="#">“About MSDP queue processing”</a> on page 235.
A user tampered with the storage.	Users must not add files to, change files on, delete files from, or change file permissions on 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 <a href="#">“About adding additional MSDP storage”</a> on page 44.
The deduplication pool is down.	Change the state to up.  See <a href="#">“Changing a Media Server Deduplication Pool state”</a> on page 224.
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 MSDP host configuration file”](#) on page 187.

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 33.

## Duplication job to PureDisk storage fails

---

**Note:** NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

---

A duplication job to PureDisk storage may fail for a large number of small files. This failure may occur if the storage unit maximum fragment is at the default setting.

Messages similar to the following may appear in the NetBackup job details:

```
2/12/2013 2:33:20 AM - Critical bpdm(pid=25204) Storage Server Error:
  (Storage server: PureDisk: msdp.) async_get_job_status:
  Replication started but failed to complete successfully: __process_batch:
  CRStoreDO failed: broken pipe, DO fingerprint d6f2521a5c66c54225ecd6a4c92.
  Look at the replication logs on the source storage server for more
  information. V-454-105
2/12/2013 2:33:22 AM - Critical bpdm(pid=25204) sts_copy_extent failed:
  error 2060014 operation aborted
2/12/2013 2:33:22 AM - end writing; write time: 3:39:53
2/12/2013 2:33:39 AM - Error bpduplicate(pid=5300) host
  msdp. backup id msdp._1360303271 optimized
  duplication failed, media write error (84).
2/12/2013 2:33:39 AM - Info msdp.(pid=25204)
  StorageServer=PureDisk:msdp.; Report=PDDO Stats for
  (msdp.): scanned: 0 KB, CR sent: 0 KB, CR sent over FC:
  0 KB, dedup: 0.0%, cache disabled
2/12/2013 2:33:40 AM - Error bpduplicate(pid=5300) Duplicate of backupid
  msdp._1360303271 failed, media write error (84).
2/12/2013 2:33:40 AM - Error bpduplicate(pid=5300) Status = no images
  were successfully processed.
2/12/2013 2:33:40 AM - end Duplicate; elapsed time: 3:41:03 media write
  error(84)
```

Messages similar to the following also may appear in a log file on the PureDisk content router that processed the duplication operation. The log file resides in a directory that corresponds to the connection to the content router for the duplication operation. The name of the directory is the NetBackup storage server IP address, and it resides in the `/Storage/log/spoold` directory on the content router.

```
spad/Store/2/266020624-1360833685.log:July 14 04:12:03 ERR [1081665856]:
25112: DO: DO 065ab494d80277148c559071d007bde1: DO record size of
378652257 bytes is larger than the maximum allowed size of 335544320
bytes: DO has too many segments
spad/Store/2/266020624-1360833685.log:February 14 04:12:03 ERR
```

```
[1081665856]: 8413608: Store terminated due to a local processing
error at: Sun Jul 14 04:12:03 2013
```

To resolve this problem, reduce the fragment size of the NetBackup storage unit that represents the PureDisk storage. Depending on the source storage content, you may have to reduce the fragment size significantly.

## MSDP no images successfully processed (191)

The `no images successfully processed (191)` message appears in the **NetBackup Administration Console Activity Monitor**.

[Table 12-5](#) describes other messages that may appear.

**Table 12-5** Case sensitivity error messages

Operation	Activity Monitor job details	Status in <code>bpdm</code> and <code>bptm</code> log files
Verify	image open failed: error 2060018: file not found	sts_open_image failed: error 2060018

The message may indicate a client name case sensitivity issue in your MSDP environment. For a description of the problem and the procedures to resolve it, see the following Symantec tech note:

<http://www.symantec.com/docs/TECH207194>

## MSDP 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 MSDP storage capacity and usage reporting”](#) on page 203.

See [“About MSDP container files”](#) on page 205.

See [“Viewing storage usage within MSDP container files”](#) on page 205.

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 MSDP queue processing”](#) on page 235.

See [“About the MSDP data removal process”](#) on page 244.

Although not advised, you can reclaim free space manually.

See “[Processing the MSDP transaction queue manually](#)” on page 235.

## Troubleshooting MSDP catalog backup

The following subsections provide information about MSDP catalog backup and recovery.

### Catalog backup

[Table 12-6](#) describes error messages that may occur when you create or update a catalog backup policy. The messages are displayed in the shell window in which you ran the `drcontrol` utility. The utility also writes the messages to its log file.

**Table 12-6** MSDP `drcontrol` codes and messages

Code or message	Description
1	Fatal error in an operating system or deduplication command that the <code>drcontrol</code> utility calls.
110	The command cannot find the necessary NetBackup configuration information.
140	The user who invoked the command does not have administrator privileges.
144	A command option or argument is required.
226	The policy name that you specified already exists.
227	This error code is passed from the NetBackup <code>bplist</code> command. The MSDP catalog backup policy you specified does not exist or no backups exist for the given policy name.
255	Fatal error in the <code>drcontrol</code> utility.

For more information about status codes and error messages, see the following:

- The Troubleshooter in the NetBackup Administration Console.
- The *NetBackup Status Codes Reference Guide* available through the following webpage:  
<http://www.symantec.com/docs/DOC5332>

### Catalog recovery from a shadow copy

If NetBackup detects corruption in the MSDP catalog, the Deduplication Manager recovers the catalog automatically from the most recent shadow copy. That recovery process also plays a transaction log so that the recovered MSDP catalog is current.

Although the shadow copy recovery process is automatic, a recovery procedure is available if you need to recover from a shadow copy manually.

See [“Restoring the MSDP catalog from a shadow copy”](#) on page 251.

## Viewing MSDP disk errors and events

You can view disk errors and events in several ways, as follows:

- The Disk Logs report.  
See [“Viewing MSDP disk reports”](#) on page 207.
- The NetBackup `bpererror` command with the `-disk` option reports on disk errors.  
The command resides in the following directories:  
UNIX: `/usr/openv/netbackup/bin/admincmd`  
Windows: `install_path\Veritas\NetBackup\bin\admincmd`

## MSDP 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 12-7 MSDP event codes and messages

Event #	Event Severity	NetBackup Severity	Message example
1000	2	Error	Operation configload/reload failed on server PureDisk:server1. on host server1..
1001	2	Error	Operation configload/reload failed on server PureDisk:server1. on host server1..
1002	4	Warning	The open file limit exceeded in server PureDisk:server1. on host server1.. Will attempt to continue further.
1003	2	Error	A connection request was denied on the server PureDisk:server1. on host server1..
1004	1	Critical	Network failure occurred in server PureDisk:server1. on host server1..
1008	2	Error	Task Aborted; An unexpected error occurred during communication with remote system in server PureDisk:server1. on host server1..

**Table 12-7** MSDP event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1009	8	Authorization	Authorization request from <IP> for user <USER> denied (<REASON>).
1010	2	Error	Task initialization on server PureDisk:server1. on host server1. got an unexpected error.
1011	16	Info	Task ended on server PureDisk:server1. on host server1..
1013	1	Critical	Task session start request on server PureDisk:server1. on host server1. got an unexpected error.
1012	2	Error	A request for agent task was denied on server PureDisk:server1. on host server1..
1014	1	Critical	Task session start request on server PureDisk:server1. on host server1. got an unexpected error.
1015	1	Critical	Task creation failed, could not initialize task class on server PureDisk:server1. on host server1..
1017	1	Critical	Service Symantec DeduplicationEngine exit on server PureDisk:server1. on host server1.. 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..
1019	1	Critical	Service Symantec DeduplicationEngine restart on server PureDisk:server1. on host server1.. 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. on host server1.. 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. on host server1.. Please check the server log for the probable cause of this error.The application has caught an unexpected signal.

**Table 12-7** MSDP event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1029	1	Critical	Double backend initialization failure; Could not initialize storage backend or cache failure detected on host PureDisk:server1. in server server1..
1030	1	Critical	Operation Storage Database Initialization failed on server PureDisk:server1. on host server1..
1031	1	Critical	Operation Content router context initialization failed on server PureDisk:server1. on host server1..
1032	1	Critical	Operation log path creation/print failed on server PureDisk:server1. on host server1..
1036	4	Warning	Operation a transaction failed on server PureDisk:server1. on host server1..
1037	4	Warning	Transaction failed on server PureDisk:server1. on host server1.. Transaction will be retried.
1040	2	Error	Operation Database recovery failed on server PureDisk:server1. on host server1..
1043	2	Error	Operation Storage recovery failed on server PureDisk:server1. on host server1..
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.
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.  The backup ID and policy name appear in the NetBackup Disk Logs report and the <code>storage_path/log/spoold/storaged.log</code> file on the storage server.
2000		Error	Low space threshold exceeded on the partition containing the storage database on server PureDisk:server1. on host server1..

# Migrating to MSDP storage

This appendix includes the following topics:

- [Migrating from PureDisk to the NetBackup MSDP](#)
- [Migrating from another storage type to MSDP](#)

## Migrating from PureDisk to the NetBackup MSDP

---

**Note:** NetBackup supports **PureDisk Deduplication Pool** storage on NetBackup 5000 series appliances only.

---

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 replace a NetBackup 5000 series appliance with the NetBackup Media Server Deduplication Option, migrate the storage to a NetBackup **Media Server Deduplication Pool**.

**Table A-1** To migrate from PureDisk to NetBackup MSDP

Step	Task	Procedure
Step 1	Install and configure NetBackup	See the <i>NetBackup Installation Guide for UNIX and Windows</i> : <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>
Step 2	Configure NetBackup deduplication	See “ <a href="#">Configuring MSDP server-side deduplication</a> ” on page 49.



**Table A-1** To migrate from PureDisk to NetBackup MSDP (*continued*)

Step	Task	Procedure
Step 3	Redirect your backup jobs	Redirect your backup jobs to the NetBackup <b>Media Server Deduplication Pool</b> .  See the <i>NetBackup Administrator's Guide, Volume I</i> :  <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a>
Step 4	Decommission your appliance	Optionally, after the backup images expire on your NetBackup 5000 series appliance, decommission your NetBackup 5000 series appliance.  See your NetBackup PureDisk documentation.

See “[Migrating from another storage type to MSDP](#)” on page 305.

## Migrating from another storage type to MSDP

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. 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 A-2** Migrating to NetBackup MSDP

Step	Task	Procedure
Step 1	Configure NetBackup deduplication	See “ <a href="#">Configuring MSDP server-side deduplication</a> ” on page 49.

**Table A-2** Migrating to NetBackup MSDP (*continued*)

Step	Task	Procedure
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, Volume I</i>:  <a href="http://www.symantec.com/docs/DOC5332">http://www.symantec.com/docs/DOC5332</a></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 MSDP](#)” on page 304.

# NetBackup appliance deduplication

This appendix includes the following topics:

- [About NetBackup appliance deduplication](#)
- [About Fibre Channel to a NetBackup 5020 appliance](#)
- [Configuring 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.. 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

You can use Fibre Channel for NetBackup MSDP storage server data traffic to a Symantec NetBackup 5020 appliance. The appliance must be at software release 1.3 or later.

Requirements for the NetBackup MSDP storage server:

- NetBackup 7.5 or later.
- A supported operating system.  
See the NetBackup operating system compatibility list on the following website.  
<http://www.netbackup.com/compatibility>
- One QLogic 2562 (ISP 2532) HBA with at least two ports.

Limitations for this solution:

- The appliance supports a maximum of 200 concurrent backup jobs.
- If no Fibre Channel connection is available, backup data travels over the IP network.

Command traffic travels over the IP network.

For each job, the job details show the amount of data that is transferred over Fibre Channel.

See “[Viewing MSDP job details](#)” on page 201.

For information about configuring the appliance and zoning the Fibre Channel SAN, see the appliance documentation.

## Configuring Fibre Channel to a NetBackup 5020 appliance

You can configure Fibre Channel for backup data or optimized duplication from a NetBackup MSDP storage server to a NetBackup 5020 appliance.

Before you do so, 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 308.

### To configure Fibre Channel to a NetBackup 5020 appliance

- 1 As the `root` user on the NetBackup media server, run the `dedup_fcmanager.sh` script with the `-e` option as in the following example:

```
/usr/openv/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]
```

- 2 Type `y` and then press Enter.

The script sets the `FIBRECHANNEL` parameter in the `pd.conf` file to 1.

See [“MSDP pd.conf file parameters”](#) on page 171.

## Disabling Fibre Channel to a NetBackup 5020 appliance

You can disable Fibre Channel from a NetBackup MSDP storage server to a NetBackup 5020 appliance.

See [“About Fibre Channel to a NetBackup 5020 appliance”](#) on page 308.

### To disable Fibre Channel to a NetBackup 5020 appliance

- 1 As the `root` user on the NetBackup media server, 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]
```

- 2 Type `y` and then press Enter.

The script sets the `FIBRECHANNEL` parameter in the `pd.conf` file to 0.

See [“MSDP pd.conf file parameters”](#) on page 171.

## 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 311.
- Statistics.  
See [“To display Fibre Channel statistics on a NetBackup 5000 series appliance”](#) on page 312.

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 308.

### To display Fibre Channel port information on a NetBackup 5020 appliance

- ◆ As the `root` user on the NetBackup media server, 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.)
/dev/sg5	SYMANTEC	192.168.0.3 (5020-Silver.)

```
**** 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 `qla2xxx` drivers do not provide KB/s output.



# Index

## A

- about NetBackup deduplication 14
- about NetBackup deduplication options 14
- about the deduplication host configuration file 187
- Active secondary operations 160
- Add at least one index marker 15
- AIR. *See* Auto Image Replication
- attributes
  - clearing deduplication pool 231
  - clearing deduplication storage server 215
  - OptimizedImage 30
  - setting deduplication pool 225
  - setting deduplication storage server 213
  - viewing deduplication pool 224
  - viewing deduplication storage server 212
- Auto Image Replication
  - about 136
  - and trusted master servers 144
  - Backup operation in source domain 134, 137
  - configuring MSDP replication to a different domain 135
  - for MSDP 146
  - overview 136
  - supported versions 137
  - synchronizing clocks of master servers 137
  - Targeted 144–145
  - topology of storage 141

## B

- Backup
  - operation 159
- backup
  - client deduplication process 269
- big endian 255, 258
- bpstinfo command 142
- byte order 255, 258

## C

- cache hits field of the job details 202

- capacity
  - adding storage 44
- capacity and usage reporting for deduplication 203
- case sensitivity
  - in NetBackup names 19
- catalog, MSDP. *See* about recovering the MSDP
- catalog. *See* MSDP catalog
- changing deduplication server hostname 216
- changing the deduplication storage server name and path 216
- CIFS 42
- clearing deduplication pool attributes 231
- client deduplication
  - components 269
  - disabling for a specific client 107
  - host requirements 26
- clustering
  - master server inter-node authentication 145
- Common Internet File System 42
- compacting container files 205
- compression
  - for MSDP backups 109
  - for MSDP optimized duplication and replication 109
  - pd.conf file setting 172
- configuring a deduplication pool 93
- configuring a deduplication storage server 67, 81
- configuring a deduplication storage unit 102
- configuring deduplication 49, 51
- container files
  - about 205
  - compaction 205
  - viewing capacity within 205
- contentrouter.cfg file
  - about 182
  - parameters for data integrity checking 240
  - ServerOptions for encryption 114
- Coordinated Universal Time (UTC) 137
- CR sent field of the job details 202–203
- credentials 28
  - adding NetBackup Deduplication Engine 222
  - changing NetBackup Deduplication Engine 222

## D

- data classifications
  - in storage lifecycle policies 153, 159
  - use of Any 154, 209
- data integrity checking
  - about deduplication 236
  - configuration settings for deduplication 239
  - configuring behavior for deduplication 237
- data removal process
  - for deduplication 244
- database system error 286
- deactivating media server deduplication 261
- dedup field of the job details 203
- deduplication
  - about credentials 28
  - about fingerprinting 60
  - about the license key 45
  - adding credentials 222
  - and Fibre Channel 42
  - and iSCSI 42
  - cache hits field of the job details 202
  - capacity and usage reporting 203
  - changing credentials 222
  - client backup process 269
  - configuration file 169
  - configuring 49, 51
  - configuring optimized synthetic backups 115
  - container files 205
  - CR sent field of the job details 202–203
  - data removal process 244
  - dedup field of the job details 203
  - event codes 301
  - license key for 45
  - licensing 46
  - limitations 24
  - media server process 266
  - network interface 29
  - node 22
  - performance 31
  - planning deployment 17
  - requirements for optimized within the same domain 119
  - scaling 33
  - scanned field of the job details 203
  - storage capacity 40
  - storage management 44
  - storage requirements 41
  - storage unit properties 103
  - deduplication configuration file
    - editing 59, 170
    - parameters 171
  - deduplication data integrity checking
    - about 236
    - configuring behavior for 237
    - settings 239
  - deduplication deduplication pool. *See* deduplication pool
  - deduplication disk volume
    - changing the state 232
    - determining the state 232
  - deduplication encryption
    - enabling for MSDP backups 112
  - deduplication host configuration file
    - about 187
    - deleting 187
  - deduplication hosts
    - and firewalls 29
    - client requirements 26
    - load balancing server 22
    - server requirements 22
    - storage server 21
  - deduplication installation
    - log file 282
  - deduplication logs
    - about 279
    - client deduplication proxy plug-in log 281
    - client deduplication proxy server log 281
    - configuration script 281
    - deduplication plug-in log 281
    - NetBackup Deduplication Engine 283
    - NetBackup Deduplication Manager 283
    - VxUL deduplication logs 279
  - deduplication node
    - about 22
    - adding a load balancing server 168
    - removing a load balancing server 218
  - deduplication optimized synthetic backups
    - about 29
  - deduplication plug-in
    - about 264
    - log file 281
  - deduplication plug-in configuration file
    - configuring 54
  - deduplication pool. *See* deduplication pool
    - about 91
    - changing properties 226
    - changing the state 224

- deduplication pool *(continued)*
    - clearing attributes 231
    - configuring 93
    - deleting 234
    - determining the state 224
    - properties 100
    - setting attributes 225
    - viewing 223
    - viewing attributes 224
  - deduplication port usage
    - about 29
    - troubleshooting 297
  - deduplication processes do not start 289
  - deduplication rate
    - how file size affects 32
    - monitoring 199
  - deduplication reference database
    - about 264
  - deduplication registry
    - resetting 187
  - deduplication servers
    - components 263
    - host requirements 22
  - deduplication storage capacity
    - about 40
    - viewing capacity in container files 205
  - deduplication storage requirements 41
  - deduplication storage server
    - about 21
    - change the name 216
    - changing properties 214
    - clearing attributes 215
    - components 263
    - configuration failure 287
    - configuring 67, 81
    - defining target for Auto Image Replication 141
    - deleting 219
    - deleting the configuration 220
    - determining the state 211
    - editing configuration file 184
    - getting the configuration 184
    - recovery 254
    - replacing the host 258
    - setting attributes 213
    - setting the configuration 186
    - viewing 211
    - viewing attributes 212
  - deduplication storage server configuration file
    - about 183
  - deduplication storage server name
    - changing 216
  - Deduplication storage unit
    - Only use the following media servers 104
    - Use any available media server 104
  - deleting backup images 234
  - deleting deduplication host configuration file 187
  - disaster recovery
    - protecting the data 37
    - recovering the storage server after catalog recovery 257
  - disk failure
    - deduplication storage server 253
  - disk logs 207
  - disk logs report 205
  - disk pool
    - cannot delete 293
  - disk pool status report 204, 207
  - disk storage unit report 207
  - Disk type 104
  - disk volume
    - changing the state 232
    - determining the state of a deduplication 232
    - volume state changes to down 291
  - domains
    - replicating backups to another. See Auto Image Replication
  - Duplication
    - optimized 160
  - Duplication to remote master. See Auto Image Replication
- ## E
- Enable file recovery from VM backup 161
  - encryption
    - enabling for MSDP backups 112
    - pd.conf file setting 173
  - endian
    - big 255, 258
    - little 255, 258
  - event codes
    - deduplication 301
- ## F
- Fibre Channel
    - and iSCSI comparison 42
    - support for 41
  - Fibre Channel and iSCSI comparison for MSDP 42

- file system
  - CIFS 42
  - NFS 42
  - Veritas File System for deduplication storage 44
- fingerprinting
  - about deduplication 60
- firewalls and deduplication hosts 29
- FlashBackup policy
  - Maximum fragment size (storage unit setting) 104
- FQDN or IP Address property in Resilient Network
  - host properties 164

## G

garbage collection. *See* queue processing

## H

host requirements 22

## I

- images on disk report 207
- Import
  - operation 154
- iSCSI
  - and Fibre Channel comparison 42
  - support for 41

## J

job ID search in unified logs 277

## L

- legacy logging 277
  - directories 278
  - locations 278
- license information failure
  - for deduplication 287
- license key
  - for deduplication 45
- licensing deduplication 46
- limitations
  - media server deduplication 24
- little endian 255, 258
- load balancing server
  - about 22
  - adding to a deduplication node 168
  - for deduplication 22
  - removing from deduplication node 218

logging

*see* legacy logging 277

logs

- about deduplication 279
- Auto Image Replication 284
- client deduplication proxy plug-in log 281
- client deduplication proxy server log 281
- deduplication configuration script log 281
- deduplication installation 282
- deduplication plug-in log 281
- disk 207
- NetBackup Deduplication Engine log 283
- NetBackup Deduplication Manager log 283
- optimized duplication 284
- VxUL deduplication logs 279

## M

- maintenance processing. *See* queue processing
- master servers
  - inter-node authentication for clustering 145
- Maximum concurrent jobs 105
- Maximum fragment size 104
- media server deduplication
  - process 266
- Media Server Deduplication Pool 91, 146
- media server deduplication pool. *See* deduplication pool
- migrating from PureDisk to NetBackup
  - deduplication 304
- migrating to NetBackup deduplication 305
- mklogdir.bat 278
- MSDP
  - replication 133
  - replication target, configuring 146
- MSDP catalog 188, 250
  - See also* MSDP catalog backup
  - See also* MSDP catalog recovery
  - about the catalog backup policy 189
  - about the shadow catalog 188
  - changing the number of shadow copies 193
  - changing the shadow catalog path 191
  - changing the shadow catalog schedule 192
  - shadow copy log files 280
- MSDP catalog backup
  - about protecting the MSDP catalog 189
  - configuring 194, 198
- MSDP catalog recovery
  - about 250
  - error codes 300

- MSDP catalog recovery *(continued)*
  - process the transaction queue. 235
  - recover from a shadow copy 251
- MSDP drcontrol utility
  - options 195
- MSDP replication
  - about 31
- MSDP storage rebasing. *See* rebasing
- mtstrm.conf file
  - configuring 54

## N

- nbstserv process 155
- NDMP
  - storage units 160
- NetBackup
  - naming conventions 19
- NetBackup deduplication
  - about 14
  - license key for 45
- NetBackup Deduplication Engine
  - about 264
  - about credentials 28
  - adding credentials 222
  - changing credentials 222
  - logs 283
- NetBackup Deduplication Manager
  - about 264
  - logs 283–284
- NetBackup deduplication options 14
- network interface
  - for deduplication 29
- NFS 42
- node
  - deduplication 22

## O

- OpenStorage
  - optimized duplication 160
- OpenStorage Disk Option 152
- optimized deduplication
  - configuring bandwidth 151
  - configuring for MSDP 126
  - limitations 119
  - logs 284
  - pull configuration within the same domain 124
  - separate network for 117

- optimized deduplication copy
  - guidance for 121
- optimized duplication
  - about 31
  - about the media server in common within the same domain 121
  - push configuration within the same domain 121
- optimized duplication encryption
  - configuring for MSDP 114
- optimized MSDP deduplication
  - requirements 119
  - within the same domain 118
- optimized synthetic backups
  - configuring for deduplication 115
  - deduplication 29
- OptimizedImage attribute 30

## P

- pd.conf file
  - about 169
  - editing 59, 170
  - parameters 171
- pdde-config.log 281
- performance
  - deduplication 31
  - monitoring deduplication rate 199
- policies
  - changing properties 161
  - creating 161
- port usage
  - and deduplication 29
  - troubleshooting 297
- Postponed secondary operations 160
- Priority for secondary operations 159
- provisioning the deduplication storage 39
- PureDisk Deduplication Pool 91

## Q

- queue processing 235
  - invoke manually 235

## R

- rebasing
  - about 241
  - FP\_CACHE\_PERIOD\_REBASING\_THRESHOLD
    - parameter 176
  - FP\_CACHE\_REBASING\_THRESHOLD
    - parameter 176

- rebasing (*continued*)
    - RebaseMaxPercentage parameter 244
    - RebaseMaxTime parameter 244
    - RebaseMinContainers parameter 244
    - RebaseScatterThreshold parameter 244
    - server-side rebasing parameters 243
  - recovery
    - deduplication storage server 254
    - from deduplication storage server disk failure 253
  - Red Hat Linux
    - deduplication processes do not start 289
  - replacing the deduplication storage server 258
  - replication
    - between NetBackup domains. *See* Auto Image Replication
    - configuring MSDP replication to a different domain 135
    - for MSDP 31, 133
    - to an alternate NetBackup domain. *See* Auto Image Replication
  - Replication Director
    - Policy Configuration Wizard, unsupported 161
  - replication encryption
    - configuring for MSDP 114
  - reports
    - disk logs 205, 207
    - disk pool status 204, 207
    - disk storage unit 207
  - reqlib directory 278
  - resetting the deduplication registry 187
  - Resiliency property in Resilient Network host properties 164
  - Resilient connection
    - Resilient Network host properties 162
  - resilient network connection
    - log file 284
  - Resilient Network host properties 162
    - FQDN or IP Address property in 164
    - Resiliency property in 164
  - restores
    - at a remote site 247
    - how deduplication restores work 246
    - specifying the restore server 248
  - reverse host name lookup
    - prohibiting 287
  - reverse name lookup 287
  - scanned field of the job details 203
  - server
    - NetBackup debug logs 278
  - server not found error 286
  - setting deduplication pool attributes 225
  - SLP Parameters 132
  - snapshots
    - operation type 159
  - spa.cfg file
    - parameters for data integrity checking 241
  - Status Collection Daemon 278
  - storage capacity
    - about 40
    - adding 44
    - for deduplication 40
    - viewing capacity in container files 205
  - Storage Lifecycle Manager service (nbstserv) 155
  - storage lifecycle policies
    - Active secondary operations 160
    - best practices document 153
    - Data classification setting 159
    - hierarchy 157
    - operations 155
    - Postponed secondary operations 160
    - Priority for secondary operations 159
    - Storage lifecycle policy name 159
    - utility 152
    - Validate Across Backup Policies button 160
  - storage paths
    - about reconfiguring 216
    - changing 216
  - storage rebasing. *See* rebasing
  - storage requirements
    - for deduplication 41
  - storage server
    - about the configuration file 183
    - change the name 216
    - changing properties for deduplication 214
    - changing the name 216
    - components for deduplication 263
    - configuring for deduplication 67, 81
    - deduplication 21
    - define target for Auto Image Replication 141
    - deleting a deduplication 219
    - deleting the deduplication configuration 220
    - determining the state of a deduplication 211
    - editing deduplication configuration file 184
    - getting deduplication configuration 184
    - recovery 254
- S**
- scaling deduplication 33

- storage server *(continued)*
  - replacing the deduplication host 258
  - setting the deduplication configuration 186
  - viewing 211
  - viewing attributes 212
- storage server configuration
  - getting 184
  - setting 186
- storage server configuration file
  - editing 184
- storage unit
  - configuring for deduplication 102
  - properties for deduplication 103
  - recommendations for deduplication 105
- storage unit groups
  - and storage lifecycle policies 160
  - not supported for Auto Image Replication
    - source 134, 137
- Storage unit name 104
- Storage unit type 104
- stream handlers
  - NetBackup 32
- synthetic backups
  - no Auto Image Replication support 137

## T

- topology of storage 141–142
- troubleshooting
  - database system error 286
  - deduplication backup jobs fail 289
  - deduplication processes do not start 289
  - general operational problems 292
  - host name lookup 287
  - installation fails on Linux 284
  - no volume appears in disk pool wizard 288
  - server not found error 286
- trusted master servers
  - for Auto Image Replication 144

## U

- unified logging 274
  - format of files 276
  - location 274
- uninstalling media server deduplication 261
- UTC (Coordinated Universal Time) 137

## V

- Validate Across Backup Policies button in SLP 160

- Validation Report tab 160
- viewing deduplication pool attributes 224
- viewing storage server attributes 212
- VM backup 161
- vmscd 278
- volume manager
  - Veritas Volume Manager for deduplication
    - storage 44
- vxlogview command 275
  - with job ID option 277

## W

- wizards
  - Policy Configuration 161