

Symantec NetBackup™ Deduplication Guide

UNIX, Windows, Linux

Release 7.6



Symantec NetBackup™ Deduplication Guide

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

Documentation version: 7.6

Legal Notice

Copyright © 2013 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, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Technical Support

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

Symantec's support offerings include the following:

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

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

www.symantec.com/business/support/

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

Contacting Technical Support

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

www.symantec.com/business/support/

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

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

North America and Latin America supportsolutions@symantec.com

Contents

Technical Support	4	
Chapter 1	Introducing the NetBackup media server deduplication option	14
	About the NetBackup deduplication options	14
	New MSDP features and enhancements for NetBackup 7.6	16
	About upgrading to NetBackup 7.6	18
Chapter 2	Planning your deployment	19
	Planning your MSDP deployment	20
	NetBackup naming conventions	21
	About the NetBackup deduplication storage destinations	22
	About the NetBackup Media Server Deduplication Option	23
	About MSDP storage servers	25
	About MSDP load balancing servers	26
	About MSDP nodes	26
	About MSDP server requirements	26
	About MSDP unsupported configurations	27
	About NetBackup MSDP Client Deduplication	28
	About MSDP client deduplication requirements and limitations	30
	About MSDP remote office client deduplication	30
	About MSDP remote client data security	31
	About remote client backup scheduling	31
	About NetBackup Deduplication Engine credentials	31
	About the network interface for MSDP	32
	About MSDP port usage	33
	About MSDP compression	33
	About MSDP encryption	34
	About MSDP optimized synthetic backups	35
	About MSDP and SAN Client	35
	About MSDP optimized duplication and replication	36
	About MSDP performance	36
	How file size may affect the MSDP deduplication rate	37
	About MSDP stream handlers	37

	MSDP deployment best practices	38
	Use fully qualified domain names	38
	About scaling MSDP	38
	Send initial full backups to the storage server	39
	Increase the number of MSDP jobs gradually	39
	Introduce MSDP load balancing servers gradually	39
	Implement MSDP client deduplication gradually	40
	Use MSDP compression and encryption	40
	About the optimal number of backup streams for MSDP	41
	About storage unit groups for MSDP	41
	About protecting the MSDP data	42
	Save the MSDP storage server configuration	42
	Plan for disk write caching	43
	How MSDP restores work	43
Chapter 3	Provisioning the storage	45
	About provisioning the MSDP storage	45
	About MSDP storage and connectivity requirements	46
	Fibre Channel and iSCSI comparison for MSDP	48
	About support for more than 32-TB of MSDP storage	49
	About MSDP storage capacity	50
	About the MSDP storage paths	50
	Do not modify MSDP storage directories and files	51
	About adding additional MSDP storage	51
	About volume management for NetBackup MSDP	51
Chapter 4	Licensing deduplication	53
	About licensing MSDP	53
	About the MSDP license key	54
	Licensing NetBackup MSDP	54
Chapter 5	Configuring deduplication	55
	Configuring MSDP deduplication	56
	Configuring MSDP client-side deduplication	59
	About the MSDP Deduplication Multi-Threaded Agent	60
	Configuring the Deduplication Multi-Threaded Agent behavior	61
	MSDP mtstrm.conf file parameters	62
	Configuring deduplication plug-in interaction with the Multi-Threaded Agent	66
	About MSDP fingerprinting	67
	About the MSDP fingerprint cache	67

Configuring the MSDP fingerprint cache behavior	68
MSDP fingerprint cache behavior options	68
About seeding the MSDP fingerprint cache for remote client deduplication	69
Seeding the MSDP fingerprint cache from the client	70
Seeding the MSDP fingerprint cache from the storage server	71
Configuring a storage server for a Media Server Deduplication Pool	72
MSDP storage paths and interface properties	77
Configuring a storage server for a PureDisk Deduplication Pool	79
About NetBackup Media Server Deduplication Pools	83
Configuring a Media Server Deduplication Pool	84
Media Server Deduplication Pool properties	91
Configuring an MSDP storage unit	93
MSDP storage unit properties	95
MSDP storage unit recommendations	96
Enabling MSDP encryption	97
Configuring optimized synthetic backups for MSDP	99
Configuring MSDP optimized duplication within the same NetBackup domain	99
About MSDP optimized duplication within the same domain	103
Configuring a separate network path for MSDP optimized duplication	115
Configuring NetBackup optimized duplication behavior	117
About configuring MSDP optimized duplication and replication bandwidth	119
Configuring MSDP replication to a different NetBackup domain	120
About MSDP replication to a different domain	122
About NetBackup Auto Image Replication	122
About trusted master servers for Auto Image Replication	130
Enabling NetBackup clustered master server inter-node authentication	131
Configuring a target for MSDP replication to a remote domain	132
About storage lifecycle policies	137
About the storage lifecycle policies required for Auto Image Replication	138
Creating a storage lifecycle policy	140
Storage Lifecycle Policy dialog box settings	143
About MSDP backup policy configuration	145
Creating a backup policy	146
Enabling MSDP client-side deduplication	147
Configuring MSDP client-direct restores	148

	Disabling MSDP client-side deduplication for a client	149
	Resilient Network properties	149
	Resilient connection resource usage	152
	Specifying resilient connections	153
	Adding an MSDP load balancing server	155
	About the MSDP pd.conf configuration file	156
	Editing the MSDP pd.conf file	156
	MSDP pd.conf file parameters	157
	About the MSDP contentrouter.cfg file	170
	About saving the MSDP storage server configuration	171
	Saving the MSDP storage server configuration	172
	Editing an MSDP storage server configuration file	172
	Setting the MSDP storage server configuration	174
	About the MSDP host configuration file	175
	Deleting an MSDP host configuration file	175
	Resetting the MSDP registry	175
Chapter 6	Migrating to MSDP storage	177
	Replacing the PureDisk Deduplication Option with MSDP on the same host	177
	Migrating from PureDisk to the NetBackup MSDP	178
	Migrating from another storage type to MSDP	179
Chapter 7	Monitoring deduplication activity	181
	Monitoring the MSDP deduplication rate	181
	Viewing MSDP job details	182
	MSDP job details	183
	About MSDP storage capacity and usage reporting	185
	About MSDP container files	187
	Viewing storage usage within MSDP container files	187
	Viewing MSDP disk reports	189
	About monitoring MSDP processes	190
	Reporting on Auto Image Replication jobs	190
Chapter 8	Managing deduplication	192
	Managing MSDP servers	192
	Viewing MSDP storage servers	193
	Determining the MSDP storage server state	193
	Viewing MSDP storage server attributes	194
	Setting MSDP storage server attributes	195
	Changing MSDP storage server properties	196

Clearing MSDP storage server attributes	197
About changing the MSDP storage server name or storage path	
path	198
Changing the MSDP storage server name or storage path	198
Removing an MSDP load balancing server	199
Deleting an MSDP storage server	201
Deleting the MSDP storage server configuration	201
About shared memory on Windows MSDP storage servers	202
Managing NetBackup Deduplication Engine credentials	202
Determining which media servers have deduplication credentials	203
Adding NetBackup Deduplication Engine credentials	203
Changing NetBackup Deduplication Engine credentials	204
Deleting credentials from a load balancing server	204
Managing Media Server Deduplication Pools	204
Viewing Media Server Deduplication Pools	205
Determining the Media Server Deduplication Pool state	205
Changing a Media Server Deduplication Pool state	205
Viewing Media Server Deduplication Pool attributes	206
Setting a Media Server Deduplication Pool attribute	206
Changing a Media Server Deduplication Pool properties	207
Clearing a Media Server Deduplication Pool attribute	212
Determining the MSDP disk volume state	213
Changing the MSDP disk volume state	213
Inventorying a NetBackup disk pool	214
Deleting a Media Server Deduplication Pool	215
Deleting backup images	215
About MSDP queue processing	216
Processing the MSDP transaction queue manually	216
About MSDP data integrity checking	217
Configuring MSDP data integrity checking behavior	218
MSDP data integrity checking configuration parameters	220
About managing MSDP storage read performance	222
About MSDP storage rebasing	223
MSDP server-side rebasing parameters	224
About the MSDP data removal process	225
Resizing the MSDP storage partition	226
About restoring files at a remote site	226
About restoring from a backup at a target master domain	227
Specifying the restore server	227

Chapter 9	Host replacement, recovery, and uninstallation	229
	Replacing the MSDP storage server host computer	229
	Recovering from an MSDP storage server disk failure	231
	Recovering from an MSDP storage server failure	233
	Recovering the MSDP storage server after NetBackup catalog recovery	235
	About uninstalling MSDP	235
	Removing MSDP	235
Chapter 10	Deduplication architecture	237
	MSDP server components	237
	Media server deduplication backup process	240
	MSDP client components	243
	MSDP client-side deduplication backup process	243
Chapter 11	Troubleshooting	248
	About unified logging	248
	About using the vxlogview command to view unified logs	249
	Examples of using vxlogview to view unified logs	250
	About legacy logging	251
	Creating NetBackup log file directories	252
	About MSDP log files	253
	Troubleshooting MSDP installation issues	258
	MSDP installation on SUSE Linux fails	258
	Troubleshooting MSDP configuration issues	259
	MSDP storage server configuration fails	259
	MSDP database system error (220)	259
	MSDP server not found error	260
	License information failure during MSDP configuration	260
	The disk pool wizard does not display an MSDP volume	261
	Troubleshooting MSDP operational issues	261
	Verify that the MSDP server has sufficient memory	262
	MSDP backup or duplication job fails	262
	MSDP client deduplication fails	264
	MSDP volume state changes to DOWN when volume is unmounted	265
	MSDP errors, delayed response, hangs	266
	Cannot delete an MSDP disk pool	266
	MSDP media open error (83)	267
	MSDP media write error (84)	269
	MSDP no images successfully processed (191)	272

	MSDP storage full conditions	272
	Viewing MSDP disk errors and events	273
	MSDP event codes and messages	273
Appendix A	NetBackup appliance deduplication	278
	About NetBackup appliance deduplication	278
	About Fibre Channel to a NetBackup 5020 appliance	279
	Enabling Fibre Channel to a NetBackup 5020 appliance	280
	Disabling Fibre Channel to a NetBackup 5020 appliance	281
	Displaying NetBackup 5020 appliance Fibre Channel port information	282
Index		285

Introducing the NetBackup media server deduplication option

This chapter includes the following topics:

- [About the NetBackup deduplication options](#)
- [New MSDP features and enhancements for NetBackup 7.6](#)
- [About upgrading to NetBackup 7.6](#)

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
<p>NetBackup Media Server Deduplication Option</p>	<p>The NetBackup Media Server Deduplication Option provides the following deduplication:</p> <ul style="list-style-type: none"> ■ 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 the storage and manages the deduplicated data. Usually, storage is a Media Server Deduplication Pool See “About the NetBackup Media Server Deduplication Option” on page 23. ■ Client-side deduplication With NetBackup MSDP client-side deduplication, clients deduplicate their backup data and then send it directly to the storage destination. A media server does not deduplicate the data. See “About NetBackup MSDP Client Deduplication” on page 28.
<p>NetBackup appliance deduplication</p>	<p>Symantec provides a hardware and a software solution that includes NetBackup deduplication. The NetBackup 5200 series of appliances run the SUSE Linux operating system and include NetBackup software and disk storage.</p> <p>The NetBackup appliances have their own documentation set.</p> <p>See “About NetBackup appliance deduplication” on page 278.</p>
<p>PureDisk deduplication</p>	<p>NetBackup PureDisk is a deduplication solution that provides bandwidth-optimized backups of data in remote offices.</p> <p>You use the PureDisk interfaces to install, configure, and manage the PureDisk servers, storage pools, and client backups. You do not use NetBackup to configure or manage the storage or backups.</p> <p>PureDisk has its own documentation set.</p> <p>See the NetBackup PureDisk Getting Started Guide.</p> <p>A PureDisk storage pool can be a storage destination for both the NetBackup Client Deduplication Option and the NetBackup Media Server Deduplication Option.</p>

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

Type	Description
PureDisk appliance deduplication	<p>Symantec provides a hardware and a software solution that includes PureDisk deduplication. The NetBackup 5000 series of appliances run the PDOS operating system and include PureDisk software and disk storage.</p> <p>The PureDisk appliances have their own documentation set. See “About NetBackup appliance deduplication” on page 278.</p>
Third- party vendor appliance deduplication	<p>The NetBackup OpenStorage option lets third-party vendor appliances function as disk storage for NetBackup.</p> <p>The disk appliance provides the storage and it manages the storage. A disk appliance may provide deduplication functionality. NetBackup backs up and restores client data and manages the life cycles of the data.</p> <p>See the NetBackup OpenStorage Solutions Guide for Disk.</p>

New MSDP features and enhancements for NetBackup 7.6

For the 7.6 release, new **Media Server Deduplication Pool** features focus on better performance and increased reliability. Collectively, the improvements provide the following benefits:

- Higher job success rate under various usage scenarios, including heavy workloads.
- Robust deduplication storage with predictable behaviors for various operations such as backup, restore, replication, and housekeeping.
- Less downtime and better system availability.

The following items describe the new features and improvements:

- Reduced storage server memory requirements.
See [“About MSDP server requirements”](#) on page 26.
- The new NetBackup Deduplication Multi-Threaded Agent improves backup performance for both client-side deduplication and media server deduplication. See [“About the MSDP Deduplication Multi-Threaded Agent”](#) on page 60.
- Improved single CPU core performance for architectures other than Intel.
- Restore operations no longer require a fingerprint index.

- Improved startup time and performance significantly for the NetBackup Deduplication Engine, as follows:
 - The engine loads only a small percentage of fingerprints into the fingerprint index cache. Previously, all fingerprints were loaded.
 See [“About the MSDP fingerprint cache”](#) on page 67.
 You can configure the cache loading behavior.
 See [“Configuring the MSDP fingerprint cache behavior”](#) on page 68.
 - The engine no longer loads the transaction queue when it starts.
 NetBackup now processes backup transactions during the backup, so they do not have to be written to the transaction queue.
- Replaced the PostgreSQL database application by a reference database for the deduplication fingerprint information. The following are some of the benefits that the new reference database provides:
 - Removes the single-point-of-failure inherent in a central reference management table.
 - Improves the performance of MSDP operations.
 - Increases the scalability of MSDP operations. More backup images can be processed and more jobs can run concurrently.
 - Limits any database corruption: Any corruption is limited to a small subset of the database. Database corruption does not affect backups, restores, or house keeping processes. Database corruption does not prevent the NetBackup Deduplication Engine (`spoold`) from starting.
 - Self-healing database: NetBackup detects corruption and reconstructs the affected parts of the database at the next backup.
- Improved replication performance.
 See [“About MSDP data integrity checking”](#) on page 217.
- Improved data integrity checking and healing.
 See [“About MSDP data integrity checking”](#) on page 217.
- A new storage rebasing method.
 After the upgrade, NetBackup uses the new method automatically. The rebasing candidates from NetBackup 7.5 are ignored.
 See [“About MSDP storage rebasing”](#) on page 223.
- NetBackup now can restore directly to a client that deduplicates its own data. Previously, restores traveled through a NetBackup media server.
 See [“Configuring MSDP client-direct restores”](#) on page 148.
- Remote WAN backup resiliency.

This feature is intended to make backups successful on the networks that reset idle socket connections within a given period of time. This feature, together with the NetBackup 7.5 WAN resiliency feature, improves job success rate for operations over WAN links with less reliable network connectivity. The success rate is also improved for operations over WAN links with intermediate firewalls and the appliances that aggressively reset idle socket connections.

Some of the improvements that the new features provide may become noticeable immediately after an upgrade to NetBackup 7.6. For example, because NetBackup uses storage more efficiently, you may see a decrease in the used disk space and a corresponding increase in the available space.

Alternatively, some of the performance improvements that the new features provide may not become evident immediately after an upgrade. However, all performance improvements should be observed for the second backups after the upgrade.

About upgrading to NetBackup 7.6

To upgrade to NetBackup 7.6, follow the procedures in the [NetBackup 7.6 Upgrade Guide](#), which you can find at the NetBackup 7.x Upgrade Portal at the following URL:

<http://www.symantec.com/docs/TECH74584>

The MSDP database changes in NetBackup 7.6 require that the existing database records be converted to the new format during an upgrade installation. During the upgrade, the installer detects if an existing MSDP configuration is present on a media server. It also determines if enough free space is available for the conversion, as follows:

Required free space for each **Media Server Deduplication Pool**: 12%.

If you do not have enough free space, you can attempt to free up storage space before you begin your upgrade. The [NetBackup 7.6 Upgrade Guide](#) provides information about how to upgrade. It includes information about how to free up storage space (if required) before you upgrade the media server. It also provides information about how much time a conversion may use.

Note: NetBackup Media Server Deduplication storage space cannot be determined accurately if you try to install NetBackup to a Solaris alternate root. Therefore, to upgrade the storage server you must first restart with the alternate root path as the running root

Planning your deployment

This chapter includes the following topics:

- [Planning your MSDP deployment](#)
- [NetBackup naming conventions](#)
- [About the NetBackup deduplication storage destinations](#)
- [About the NetBackup Media Server Deduplication Option](#)
- [About NetBackup MSDP Client Deduplication](#)
- [About MSDP remote office client deduplication](#)
- [About NetBackup Deduplication Engine credentials](#)
- [About the network interface for MSDP](#)
- [About MSDP port usage](#)
- [About MSDP compression](#)
- [About MSDP encryption](#)
- [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	See “About the NetBackup deduplication storage destinations” on page 22.
Step 2	Determine which type of deduplication to use	See “About the NetBackup Media Server Deduplication Option” on page 23. See “About NetBackup MSDP Client Deduplication” on page 28. See “About MSDP remote office client deduplication” on page 30.
Step 3	Determine the requirements for deduplication hosts	See “About MSDP storage servers” on page 25. See “About MSDP server requirements” on page 26. See “About MSDP client deduplication requirements and limitations” on page 30. See “About the network interface for MSDP” on page 32. See “About MSDP port usage” on page 33. See “About scaling MSDP” on page 38. See “About MSDP performance” on page 36.
Step 4	Determine the credentials for deduplication	See “About NetBackup Deduplication Engine credentials” on page 31.
Step 5	Read about compression and encryption	See “About MSDP compression” on page 33. See “About MSDP encryption” on page 34.
Step 6	Read about optimized synthetic backups	See “About MSDP optimized synthetic backups” on page 35.
Step 7	Read about deduplication and SAN Client	See “About MSDP and SAN Client” on page 35.
Step 8	Read about optimized duplication and replication	See “About MSDP optimized duplication and replication” on page 36.
Step 9	Read about stream handlers	See “About MSDP stream handlers” on page 37.
Step 10	Read about best practices for implementation	See “MSDP deployment best practices” on page 38.

Table 2-1 Deployment overview (*continued*)

Step	Deployment task	Where to find the information
Step 11	Determine the storage requirements and provision the storage	See “About provisioning the MSDP storage” on page 45. See “About MSDP storage and connectivity requirements” on page 46. See “About MSDP storage capacity” on page 50. See “About the MSDP storage paths” on page 50.
Step 12	License MSDP	See “About licensing MSDP” on page 53.
Step 13	Configure MSDP	See “Configuring MSDP deduplication” on page 56. See “Configuring MSDP client-side deduplication” on page 59.
Step 14	Migrate from other storage to NetBackup deduplication	See “Replacing the PureDisk Deduplication Option with MSDP on the same host” on page 177. See “Migrating from PureDisk to the NetBackup MSDP” on page 178. See “Migrating from another storage type to MSDP” on page 179.

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 NetBackup Deduplication Engine credentials”](#) on page 31.

About the NetBackup deduplication storage destinations

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

Table 2-2 NetBackup deduplication storage destinations

Destination	Description
Media Server Deduplication Pool	<p>A NetBackup Media Server Deduplication Pool represents the disk storage that is attached to a NetBackup media server. NetBackup deduplicates the data and hosts the storage.</p> <p>If you use this destination, use this guide to plan, implement, configure, and manage deduplication and the storage. When you configure the storage server, select Media Server Deduplication Pool as the storage type.</p> <p>For a Media Server Deduplication Pool storage destination, all hosts that are used for the deduplication must be NetBackup 7.0 or later. Hosts include the master server, the media servers, and the clients that deduplicate their own data. Integrated deduplication means that the components installed with NetBackup perform deduplication.</p> <p>A NetBackup 5200 series appliance also provides Media Server Deduplication Pool destinations to which NetBackup can send deduplicated data.</p>
PureDisk Deduplication Pool	<p>A NetBackup PureDisk Deduplication Pool represents a PureDisk storage pool. NetBackup deduplicates the data, and the PureDisk environment hosts the storage.</p> <p>If you use a PureDisk Deduplication Pool, use the following documentation:</p> <ul style="list-style-type: none"> ■ The PureDisk documentation to plan, implement, configure, and manage the PureDisk environment, which includes the storage. See the NetBackup PureDisk Getting Started Guide. ■ This guide to configure backups and deduplication in NetBackup. When you configure the storage server, select PureDisk Deduplication Pool as the storage type. <p>A PureDisk Deduplication Pool destination requires that PureDisk be at release 6.6 or later.</p> <p>A NetBackup 5000 series appliance also provides a PureDisk storage pool to which NetBackup can send deduplicated data.</p>

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

Destination	Description
PureDisk Storage Pool Authority	<p>The PureDisk Deduplication Option (PDDO) exposes PureDisk Storage Pool Authority storage to NetBackup as a storage destination. PDDO was introduced before Symantec added integrated deduplication into NetBackup in the 7.0 release.</p> <p>If you use this storage option, use the PureDisk documentation to plan, implement, configure, and manage the storage and to configure NetBackup to use the agent.</p> <p>The PureDisk Deduplication Option does not support IPV6. You should use only IPV4 addressing on NetBackup media servers that host a PureDisk Deduplication Option agent. The PureDisk Deduplication Option is an agent that is part of the PureDisk software. It is packaged, delivered, licensed, and purchased separately from NetBackup.</p> <p>See the Symantec NetBackup PureDisk™ Deduplication Option Guide that is included with the PureDisk distribution.</p>

About the NetBackup Media Server Deduplication Option

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

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

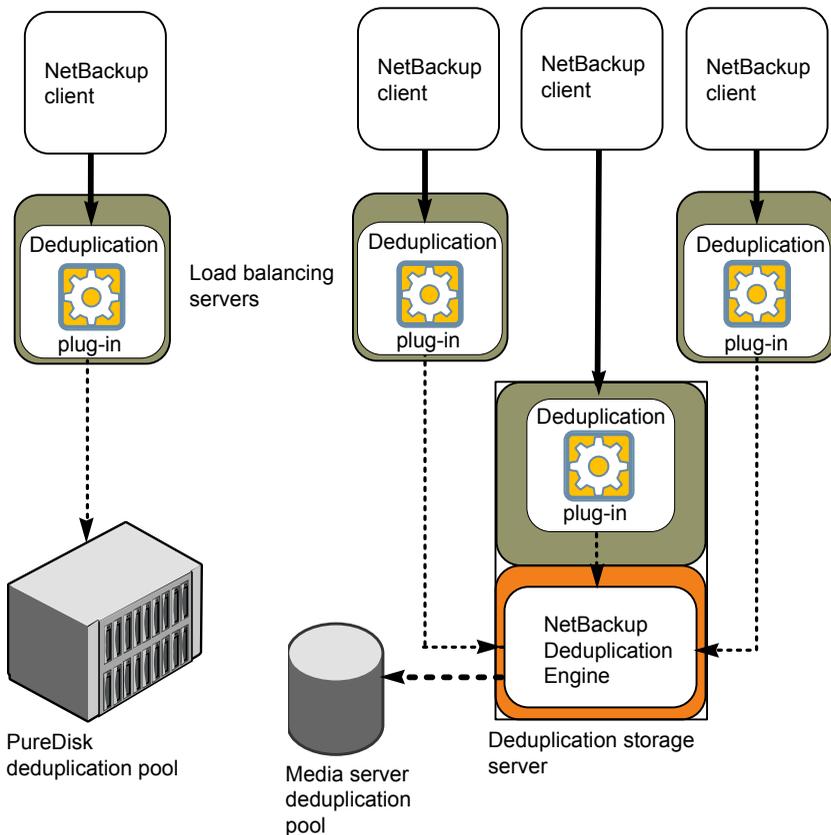
See [“About MSDP storage servers”](#) on page 25.

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

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

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

Figure 2-1 NetBackup media server deduplication



A PureDisk storage pool may also be the storage destination.

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

More detailed information is available.

See [“MSDP server components”](#) on page 237.

See [“Media server deduplication backup process”](#) on page 240.

About MSDP storage servers

One host functions as the storage server for a deduplication node, as follows:

- For a NetBackup **Media Server Deduplication Pool**, the host must be a NetBackup media server.
- For a **PureDisk Deduplication Pool**, the host must be a PureDisk host.

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

The storage server does the following:

- Writes the data to and reads data from the disk storage.
- Manages that storage.

For a **Media Server Deduplication Pool**, the storage server also deduplicates data. Therefore, one host both deduplicates the data and manages the storage.

You can configure other NetBackup media servers to deduplicate data also.

See [“About MSDP load balancing servers”](#) on page 26.

Only one storage server exists for each NetBackup deduplication node.

See [“About MSDP nodes”](#) on page 26.

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 103.

- 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 122.

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 39.

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 2-3](#) 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 2-3 MSDP server minimum requirements

Component	Storage server	Load balancing server or PureDisk Deduplication Option host
CPU	<p>Symantec recommends at least a 2.2-GHz clock rate. A 64-bit processor is required.</p> <p>At least four cores are required. Symantec recommends eight cores.</p> <p>For 64 TBs of storage, Intel x86-64 architecture requires eight cores.</p>	<p>Symantec recommends at least a 2.2-GHz clock rate. A 64-bit processor is required.</p> <p>At least two cores are required. Depending on throughput requirements, more cores may be helpful.</p>
RAM	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 NetBackup landing page on the Symantec Support Web site.</p>	<p>The operating system must be a supported 64-bit operating system.</p> <p>See the operating system compatibility list for your NetBackup release on the NetBackup landing page on the Symantec Support Web site.</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 nondeduplication storage.) If you perform more than 100 backups a day, deduplication operations may affect master server operations.

See [“About MSDP performance”](#) on page 36.

See [“About MSDP queue processing”](#) on page 216.

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 media server deduplication and the PureDisk PDDO agent cannot reside on the same host. If the NetBackup installation process detects the PDDO agent, the installation fails. To upgrade a NetBackup media server that hosts a PDDO agent, you must first remove the PDDO agent. You then can use that host as a front end for your PureDisk Storage Pool Authority. (The host must be a host type that is supported for NetBackup deduplication.)
 See [“Replacing the PureDisk Deduplication Option with MSDP on the same host”](#) on page 177.
 See the [NetBackup PureDisk Deduplication Option \(PDDO\) Guide](#).
- NetBackup does not support clustering of deduplication storage servers or load balancing servers.
- NetBackup deduplication components cannot reside on the same host as the PureDisk Deduplication Option (PDDO) agent that is installed from the PureDisk distribution.
- Deduplication within each media server deduplication node is supported; global deduplication between nodes is not supported.

About NetBackup MSDP Client Deduplication

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

With NetBackup Client Deduplication (also know as *client-side deduplication*), the client hosts the deduplication plug-in that duplicates the backup data. The NetBackup client software creates the image of backed up files as for a normal backup. Next, the deduplication plug-in breaks the backup image into segments and compares them to all of the segments that are stored in that deduplication node. The plug-in then sends only the unique segments to the NetBackup Deduplication Engine on the storage server. The engine writes the data to a **Media Server Deduplication Pool**.

Client deduplication does the following:

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

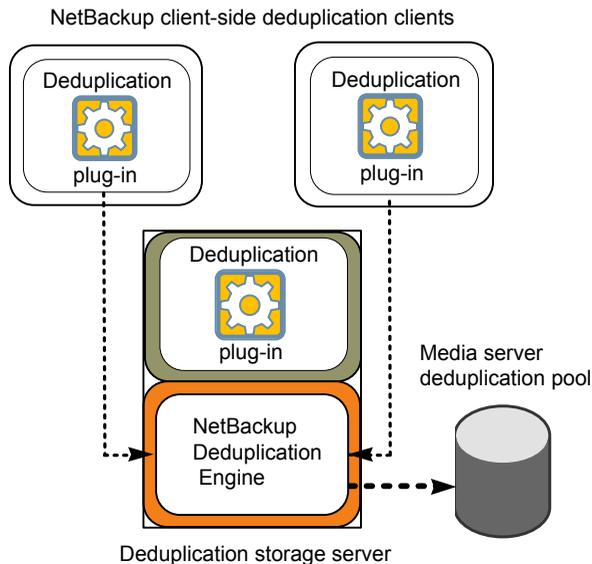
NetBackup Client Deduplication is a solution for the following cases:

- Remote office or branch office backups to the data center.
 NetBackup provides resilient network connections for remote office backups.
 See [“About MSDP remote office client deduplication”](#) on page 30.
- LAN connected file server
- Virtual machine backups.

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

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

Figure 2-2 NetBackup client deduplication



A PureDisk storage pool may also be the storage destination.

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

More information is available.

See [“About MSDP remote office client deduplication”](#) on page 30.

See [“MSDP client components”](#) on page 243.

See [“MSDP client-side deduplication backup process”](#) on page 243.

About MSDP client deduplication requirements and limitations

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.

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 Release Notes](#).

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 149.

See [“Specifying resilient connections”](#) on page 153.

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 157.

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 69.

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 34.

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

See [“How MSDP restores work”](#) on page 43.

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 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 can be up to 127 characters in length. The user name cannot be empty or all space characters.
- The password can be up to 100 characters in length. 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 198.

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 2-4 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 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 clients that deduplicate their own data.
443	The PureDisk Storage Pool Authority. Open this port between the NetBackup clients that deduplicate their own data and the PureDisk Storage Pool.

About MSDP compression

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

Symantec recommends that you use the deduplication compression.

Note: Do not enable compression by selecting the **Compression** setting on the **Attributes** tab of the **Policy** dialog box. If you do, the NetBackup compresses the data before it reaches the deduplication plug-in that deduplicates it. Consequently, deduplication rates are very low. Also, NetBackup does not use the Deduplication Multi-Threaded Agent if policy-based compression is enabled.

See [“About the MSDP Deduplication Multi-Threaded Agent”](#) on page 60.

See [“Use MSDP compression and encryption”](#) on page 40.

See [“About the MSDP pd.conf configuration file”](#) on page 156.

See [“MSDP pd.conf file parameters”](#) on page 157.

About MSDP encryption

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

Symantec recommends that you use deduplication encryption.

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

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

Deduplication uses the Blowfish algorithm for encryption.

Note: Do not enable encryption by selecting the **Encryption** setting on the **Attributes** tab of the **Policy** dialog box. If you do, NetBackup encrypts the data before it reaches the deduplication plug-in that deduplicates it. Consequently, deduplication rates are very low. Also, NetBackup does not use the Deduplication Multi-Threaded Agent if policy-based encryption is enabled.

See [“About the MSDP Deduplication Multi-Threaded Agent”](#) on page 60.

See [“Use MSDP compression and encryption”](#) on page 40.

See [“Enabling MSDP encryption”](#) on page 97.

See [“About the MSDP pd.conf configuration file”](#) on page 156.

See [“MSDP pd.conf file parameters”](#) on page 157.

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 99.

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 195.

See [“Setting a Media Server Deduplication Pool attribute”](#) on page 206.

Table 2-5 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 2-6 NetBackup OpenStorage optimized duplication and replication methods

Optimized duplication method	Description
Within the same NetBackup domain	See “About MSDP optimized duplication within the same domain” on page 103.
To a remote NetBackup domain	See “About NetBackup Auto Image Replication” on page 122.

About MSDP performance

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

[Table 2-7](#) 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 26.

Table 2-7 MSDP job load performance for an MSDP storage server

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

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 269.

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 67.

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 36.

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 36.

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 33.

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, as follows:

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

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 2-8 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 Media Server Deduplication Pool storage units and PureDisk storage pool storage units.

Table 2-8 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"> ■ 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 “About MSDP optimized duplication within the same domain” on page 103. ■ Optimized synthetic backups. If NetBackup cannot produce the optimized synthetic backup, NetBackup creates the more data-movement intensive synthetic backup. See “About MSDP optimized synthetic backups” on page 35.

About protecting the MSDP data

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

- Use NetBackup optimized duplication to copy the images to another deduplication node off-site location.
 Optimized duplication copies the primary backup data to another deduplication pool. It provides the easiest, most efficient method to copy data off-site yet remain in the same NetBackup domain. You then can recover from a disaster that destroys the storage on which the primary copies reside by retrieving images from the other deduplication pool.
 See [“About MSDP optimized duplication within the same domain”](#) on page 103.
 See [“Configuring MSDP optimized duplication within the same NetBackup domain”](#) on page 99.
- For the primary deduplication storage, use a SAN volume with resilient storage methodologies to replicate the data to a remote site. If the deduplication database is on a different SAN volume, replicate that volume to the remote site also.

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

Save the 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 171.

See [“Saving the MSDP storage server configuration”](#) on page 172.

See [“Editing an MSDP storage server configuration file”](#) on page 172.

Plan for disk write caching

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

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

Protection can include the following:

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

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

How MSDP restores work

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

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

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

- A load balancing server in the same deduplication node.

- A deduplication server in a different deduplication node that is the target of optimized duplication.
Optimized duplication requires a server in common between the two deduplication nodes.
See [“About the media servers for MSDP optimized duplication within the same domain”](#) on page 106.

You can specify the server to use for restores.

See [“Specifying the restore server”](#) on page 227.

Provisioning the storage

This chapter includes the following topics:

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

About provisioning the MSDP storage

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

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

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

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

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

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

- Storage for the backups in the originating domain. This storage contains your client backups. It is the source for the 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 122.

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

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

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

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 requires dedicated HBAs and perhaps 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.
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.

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

Item	Fibre Channel	iSCSI
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 267.

See [“MSDP media write error \(84\)”](#) on page 269.

About support for more than 32-TB of MSDP storage

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

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

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

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

- Stop the deduplication services (`spad` and `spoold`) on the storage server host.
- Grow the disk storage to the desired size up to 64 TBs.
See [“About adding additional MSDP storage”](#) on page 51.
- If more CPU and memory are required, shut down the storage server host and then upgrade the CPU and add the required amount of memory.
- Start the host.
- Upgrade NetBackup on the storage server host.

About MSDP storage capacity

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

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

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

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

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

About the MSDP storage paths

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

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

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

NetBackup requires specific characters for the path names.

See [“MSDP storage paths and interface properties”](#) on page 77.

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

See [“About MSDP storage and connectivity requirements”](#) on page 46.

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, carefully decide during the planning phase where and how you want the deduplicated backup data stored.

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 226.

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, be aware of the following:

- NetBackup supports deduplication on specific 64-bit host operating systems. If you intend to upgrade an existing media server and use it for deduplication, that host must be supported.
For the supported systems, see the [NetBackup Release Notes](#).
- NetBackup deduplication components cannot reside on the same host as a PureDisk Deduplication Option agent.
To use a PDDO agent host for NetBackup deduplication, first remove the PDDO agent from that host.
See the [NetBackup PureDisk Deduplication Option \(PDDO\) Guide](#).
Then, upgrade that host to NetBackup 7.0 or later.
Finally, configure that host as a deduplication storage server or as a load balancing server.

About the 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](#).

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 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](#)
- [Seeding the MSDP fingerprint cache from the client](#)
- [Seeding the MSDP fingerprint cache from the storage server](#)
- [Configuring a storage server for a Media Server Deduplication Pool](#)
- [Configuring a storage server for a PureDisk Deduplication Pool](#)
- [About NetBackup Media Server Deduplication Pools](#)
- [Configuring a Media Server Deduplication Pool](#)
- [Configuring an MSDP storage unit](#)
- [Enabling MSDP encryption](#)
- [Configuring optimized synthetic backups for MSDP](#)

- Configuring MSDP optimized duplication within the same NetBackup domain
- Configuring MSDP replication to a different NetBackup domain
- 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
- Enabling MSDP client-side deduplication
- Configuring MSDP client-direct restores
- Disabling MSDP client-side deduplication for a client
- 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

Configuring MSDP 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](#).

Table 5-1 MSDP configuration tasks

Step	Task	Procedure
Step 1	Install the license key	See "Licensing NetBackup MSDP" on page 54.
Step 2	Create NetBackup log file directories on the master server and the media servers	See "About MSDP log files" on page 253. See "Creating NetBackup log file directories" on page 252.
Step 3	Configure the Deduplication Multi-Threaded Agent behavior	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. See "About the MSDP Deduplication Multi-Threaded Agent" on page 60. See "Configuring the Deduplication Multi-Threaded Agent behavior" on page 61. See "Configuring deduplication plug-in interaction with the Multi-Threaded Agent" on page 66.
Step 4	Configure the fingerprint cache behavior	Configuring the fingerprint cache behavior is optional. See "About MSDP fingerprinting" on page 67. See "About the MSDP fingerprint cache" on page 67. See "Configuring the MSDP fingerprint cache behavior" on page 68.
Step 5	Configure a deduplication storage server	How many storage servers you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication. See "About MSDP storage servers" on page 25. See "About MSDP optimized duplication and replication" on page 36. Which type of storage server to configure depends on the storage destination. See "About the NetBackup deduplication storage destinations" on page 22. See "Configuring a storage server for a Media Server Deduplication Pool" on page 72. See "Configuring a storage server for a PureDisk Deduplication Pool" on page 79.

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

Step	Task	Procedure
Step 6	Configure a disk pool	How many disk pools you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication. See “About NetBackup Media Server Deduplication Pools” on page 83. See “Configuring a Media Server Deduplication Pool” on page 84.
Step 7	Configure a storage unit	See “Configuring an MSDP storage unit” on page 93.
Step 8	Enable encryption	Encryption is optional. See “Enabling MSDP encryption” on page 97.
Step 9	Configure optimized synthetic backups	Optimized synthetic backups are optional. See “Configuring optimized synthetic backups for MSDP” on page 99.
Step 10	Configure optimized duplication copy	Optimized duplication is optional. See “Configuring a separate network path for MSDP optimized duplication” on page 115. See “Configuring MSDP optimized duplication within the same NetBackup domain” on page 99.
Step 11	Configure replication	Replication is optional. See “About the replication topology for Auto Image Replication” on page 127. See “Configuring a target for MSDP replication to a remote domain” on page 132. See “Viewing the replication topology for Auto Image Replication” on page 128. See “About the storage lifecycle policies required for Auto Image Replication” on page 138. See “Creating a storage lifecycle policy” on page 140.
Step 12	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 “About MSDP backup policy configuration” on page 145. See “Creating a backup policy” on page 146.

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

Step	Task	Procedure
Step 13	Specify advanced deduplication settings	<p>Advanced settings are optional.</p> <p>See “About the MSDP pd.conf configuration file” on page 156.</p> <p>See “Editing the MSDP pd.conf file” on page 156.</p> <p>See “MSDP pd.conf file parameters” on page 157.</p>

Configuring MSDP client-side deduplication

This topic describes how to configure client deduplication in NetBackup.

Table 5-2 Client deduplication configuration tasks

Step	Task	Procedure
Step 1	Configure media server deduplication	See “Configuring MSDP deduplication” on page 56.
Step 2	Learn about client deduplication	See “About NetBackup MSDP Client Deduplication” on page 28.
Step 3	Configure a resilient connection for remote office clients	<p>Resilient connections are optional.</p> <p>See “About MSDP remote office client deduplication” on page 30.</p> <p>See “Resilient Network properties” on page 149.</p> <p>See “Specifying resilient connections” on page 153.</p>
Step 4	Enable client-side deduplication	See “Enabling MSDP client-side deduplication” on page 147.
Step 5	Configure remote client fingerprint cache seeding	<p>Configuring remote client fingerprint cache seeding is optional.</p> <p>See “Seeding the MSDP fingerprint cache from the client” on page 70.</p> <p>See “About seeding the MSDP fingerprint cache for remote client deduplication” on page 69.</p> <p>See “Seeding the MSDP fingerprint cache from the storage server” on page 71.</p>

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

Step	Task	Procedure
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 “Configuring MSDP client-direct restores” on page 148.

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 “Configuring the Deduplication Multi-Threaded Agent behavior” on page 61.
Whether or not the deduplication plug-in sends backups to the Multi-Threaded Agent.	See “Configuring deduplication plug-in interaction with the Multi-Threaded Agent” on page 66.
The clients that should use the Deduplication Multi-Threaded Agent for backups.	See “Configuring deduplication plug-in interaction with the Multi-Threaded Agent” on page 66.
The backup policies that should use the Deduplication Multi-Threaded Agent.	See “Configuring deduplication plug-in interaction with the Multi-Threaded Agent” on page 66.

[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"> ■ Virtual synthetic backups ■ NetBackup Accelerator ■ SEGKSIZE is greater than 128 (<code>pd.conf</code> file) ■ DONT_SEGMENT_TYPES enabled (<code>pd.conf</code> file) ■ MATCH_PDRO = 1 (<code>pd.conf</code> file) <p>See “MSDP pd.conf file parameters” on page 157.</p>
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 “About MSDP compression” on page 33.</p> <p>See “About MSDP encryption” on page 34.</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 60.

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 62.

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 61.

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)

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"> ■ Windows: <code>LogPath=install_path\Veritas\pdde\...\netbackup\logs\pdde</code> ■ UNIX: <code>LogPath=/var/log/puredisk</code>
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.<date/time stamp></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)

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 “MSDP <code>pd.conf</code> file parameters” on page 157.</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 Table 5-7). (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 “MSDP <code>pd.conf</code> file parameters” on page 157.</p> <p>Possible values: 1-32, inclusive.</p> <p>Warning: 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 $(\text{BackupReadBufferCount} * \text{BackupReadBufferSize}) + (3 * \text{BackupShmBufferSize}) + \text{FpCacheMaxMbSize}$ (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)

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)

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"> ■ Windows and Linux: The threshold value is 2. ■ Solaris: The threshold value is 4. <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 pd.conf file parameters” on page 157.</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 156.

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 157.
- 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 67.

See [“Media server deduplication backup process”](#) on page 240.

See [“MSDP client-side deduplication backup process”](#) on page 243.

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 67.

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 68.

You can also control the fingerprint cache seeding for clients.

See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 69.

Configuring the MSDP fingerprint cache behavior

You can configure the cache loading behavior.

See [“About the MSDP fingerprint cache”](#) on page 67.

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 68.

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 170.

Table 5-8 Cache load parameters

Behavior	Description
NumConWriteFiles	<p>The number of fingerprints to load at startup.</p> <p>The <code>NumConWriteFiles</code> controls the number of container files for which to load fingerprints. NetBackup loads fingerprints from the highest to lowest container number. That is, from the most recently created container, to the next most recently created container, and so on.</p> <p>The default value is 1024. One-thousand-twenty-four concurrent backup streams each have their own container to write into.</p>
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 “Seeding the MSDP fingerprint cache from the client” on page 70.
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 “Seeding the MSDP fingerprint cache from the storage server” on page 71.

See [“About the MSDP fingerprint cache”](#) on page 67.

Seeding the MSDP fingerprint cache from 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 69.

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 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:

clienthostmachine, backuppolicy, date

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

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 156.

See [“MSDP pd.conf file parameters”](#) on page 157.

Seeding the MSDP fingerprint cache from 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:

storage_path/databases/catalog/2/#pdseed/client_name

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 69.

To seed the fingerprint cache for the storage server

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

UNIX: `/usr/opensv/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`

Repeat the command for each client for which you want to seed with fingerprints.

- 2 You can verify that the seeding directory for the client was created by using the following command:

```
seedutil -list_clients
```

- 3 After the client or clients are backed up, remove the seeding directory for the client or clients. The following is the command syntax:

```
seedutil -clear client_name
```

NetBackup clears the seeding directory automatically when a backup image for the client expires. However, Symantec recommends that you clear the client seeding directories manually.

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 25.

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

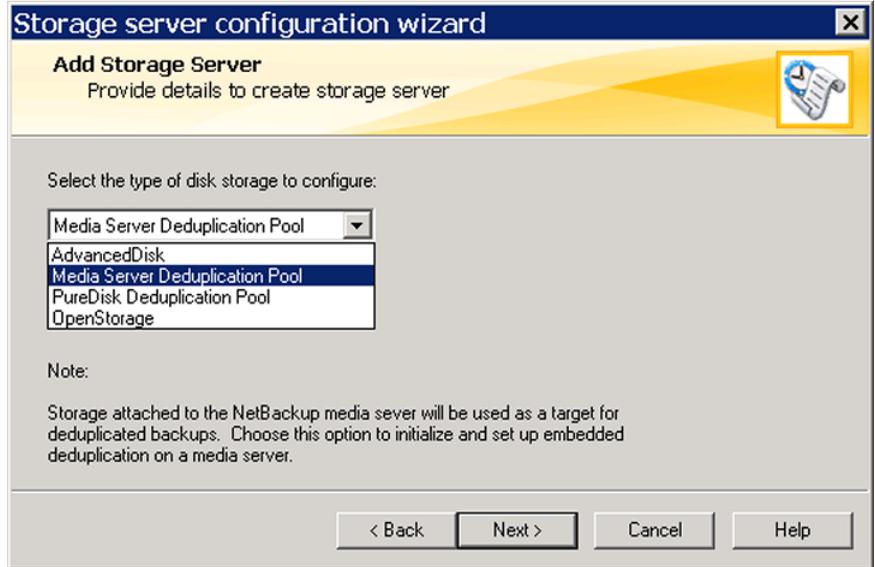
To configure a NetBackup storage server for a Media Server Deduplication Pool

- 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**.

- 3 Click **Next** on the welcome panel of the wizard.

The **Add Storage Server** panel appears.

The following is an example of the wizard panel:

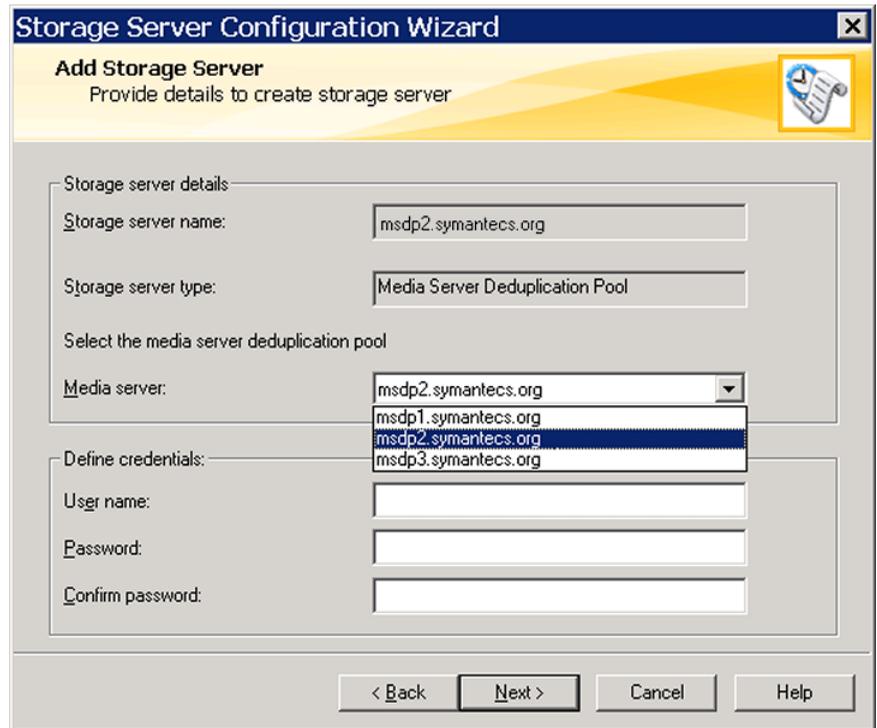


- 4 On the **Add Storage Server** panel, select **Media Server Deduplication Pool** from the drop-down menu.

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

After you select **Media Server Deduplication Pool**, click **Next**. A wizard panel about **Storage server details** appears.

The following is an example of the wizard panel:



- From the **Media server** drop-down menu in the **Storage server details** pane, select the media server that you want to configure as the storage server.

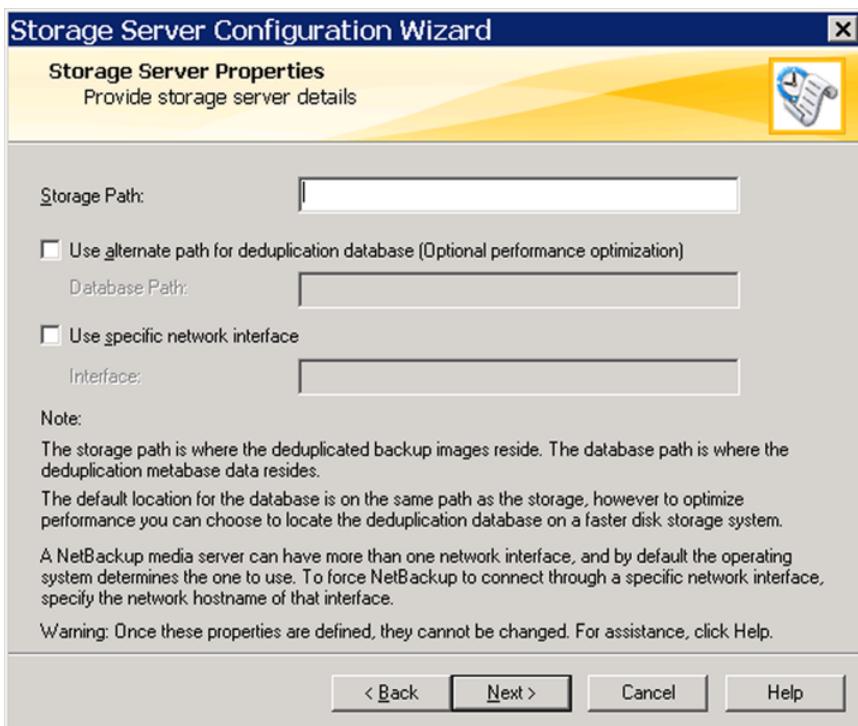
Also enter the credentials for the NetBackup Deduplication Engine. The deduplication components use the credentials when they communicate with the NetBackup Deduplication Engine. The credentials are for the engine, not for the host on which it runs.

See [“About NetBackup Deduplication Engine credentials”](#) on page 31.

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

The **Storage Server Properties** panel appears.

The following is an example of the wizard panel:

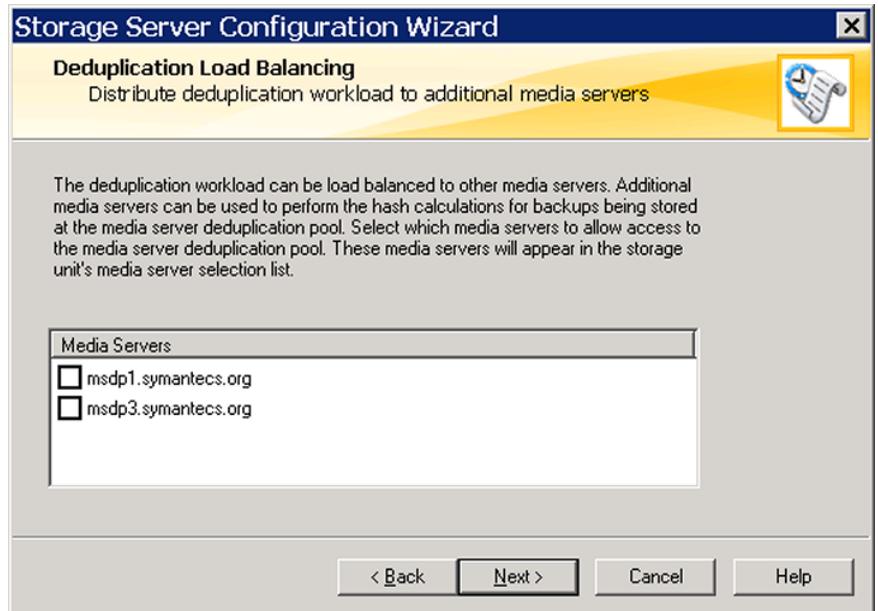


- 6 On the **Storage Server Properties** panel, enter or select the properties for the deduplication storage server.

See [“MSDP storage paths and interface properties”](#) on page 77.

After you click **Next**, the **Deduplication Load Balancing** panel appears.

The following is an example of the wizard panel:



- 7 The **Deduplication Load Balancing** panel lets you select other NetBackup media servers for deduplication.

See [“About MSDP load balancing servers”](#) on page 26.

The wizard populates the media servers with the NetBackup Deduplication Engine credentials you entered on a previous panel.

- 8 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.

- 9 After NetBackup creates the storage server, click **Next** on the **Storage Server Creation Status** panel.

If storage server creation fails, troubleshooting help is available.

See [“Troubleshooting MSDP configuration issues”](#) on page 259.

- 10 After the storage server is created, you can do the following.
- To configure a disk pool for this storage server, click **Next**.
See [“Configuring a Media Server Deduplication Pool”](#) on page 84.
 - To exit the wizard, click **Close**.

MSDP storage paths and interface properties

Media Server Deduplication Pool only.

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

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

Table 5-10 MSDP storage server properties

Property	Description
Storage path	<p>The path to the storage.</p> <p>Because the storage requires a directory path, do not use only the root node (/) or drive letter (C:\) as the storage path.</p> <p>You can use the following characters in the storage path name:</p> <ul style="list-style-type: none"> ■ 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. ■ Any integer from 0 to 9, inclusive. ■ A space character. ■ Any of the following characters: UNIX: _ - : . / \ Windows: _ - : . \ (a colon (:)) is allowed only after a drive letter (for example, C:\MSDP_Storage) <p>See “About the MSDP storage paths” on page 50.</p>

Table 5-10 MSDP storage server properties (*continued*)

Property	Description
Use alternate path for deduplication database	<p>By default, NetBackup uses the storage path for the database location.</p> <p>Select this option to use a location other than the default for the deduplication database. The database is separate from the NetBackup catalogs.</p> <p>For performance optimization, Symantec recommends that you use a separate disk, volume, partition, or spindle for the deduplication database.</p>
Database path	<p>If you selected Use alternate path for deduplication database, enter the path name for the database.</p> <p>You can use the following characters in the path name:</p> <ul style="list-style-type: none"> ■ 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. ■ Any integer from 0 to 9, inclusive. ■ A space character. ■ Any of the following characters: UNIX: _ - : . / \ Windows: _ - : . \ (a colon (:) is allowed only after a drive letter (for example, C:\MSDP_Storage)
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 “About the network interface for MSDP” on page 32.</p>
Interface	<p>If you selected Use specific network interface, 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

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

See [“About MSDP storage servers”](#) on page 25.

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

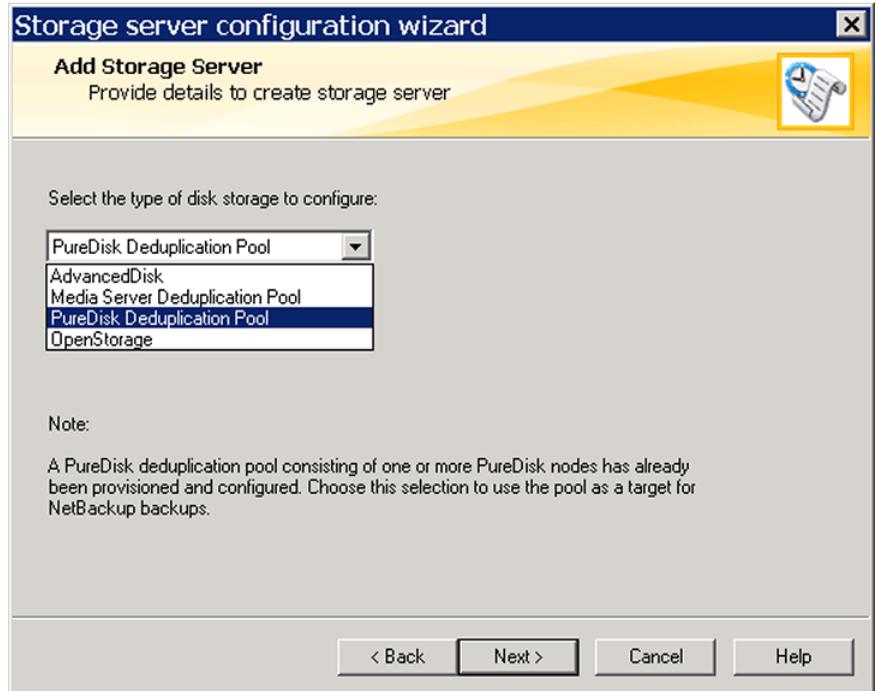
To configure a NetBackup storage server for a PureDisk Deduplication Pool by using the wizard

- 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**.

- 3 Click **Next** on the welcome panel of the wizard.

The **Add Storage Server** panel appears.

The following is an example of the wizard panel:

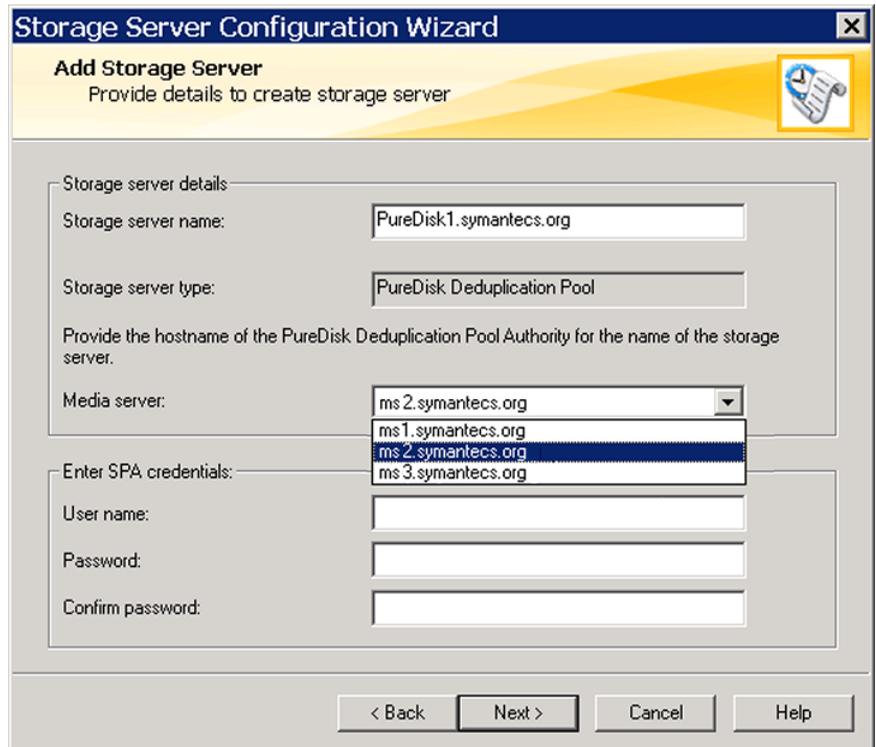


- 4 On the **Add Storage Server** panel, select **PureDisk Deduplication Pool** from the drop-down menu.

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

After you select **PureDisk Deduplication Pool**, click **Next**. A wizard panel about **Storage server details** appears.

The following is an example of the wizard panel:



- 5 Enter the host name of the PureDisk Storage Pool Authority in the **Storage server name** field.

Next, select a media server to use to query the PureDisk Storage Pool Authority host for its capabilities. That media server also is used for deduplication. You can add other deduplication load balancing servers on the next wizard panel.

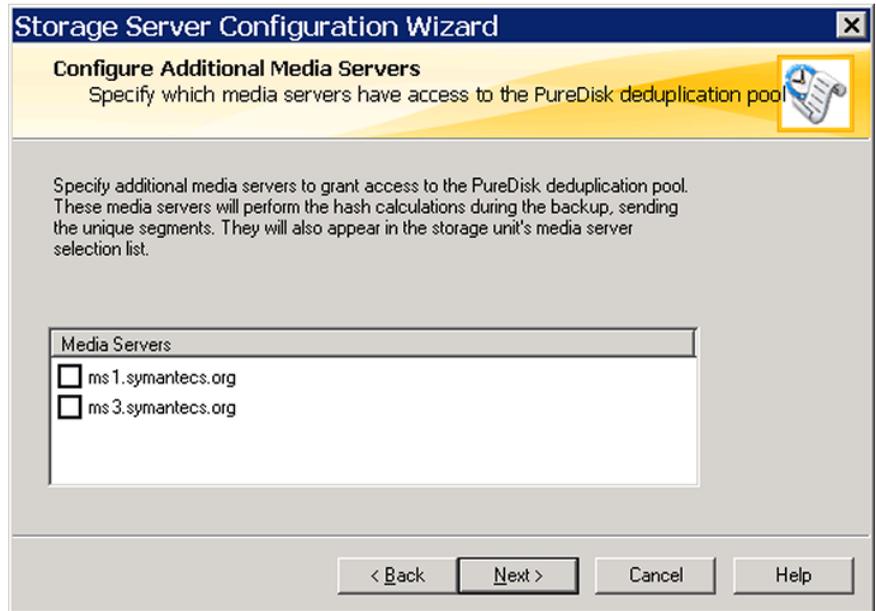
Finally, enter the credentials for the PureDisk Storage Pool Authority:

Username	The user name for the Storage Pool Authority.
Password	The password for the Storage Pool Authority.
Confirm password	To confirm the password, re-enter the password.

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

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

The following is an example of the wizard panel:



- 6 The **Configure Additional Media Servers** panel lets you select other NetBackup media servers for deduplication.

See “[About MSDP load balancing servers](#)” on page 26.

The wizard populates the media servers with the Storage Pool Authority credentials you entered on a previous panel.
- 7 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.
- 8 After NetBackup creates the storage server, click **Next** on the **Storage Server Creation Status** panel.

If storage server creation fails, troubleshooting help is available.

See “[Troubleshooting MSDP configuration issues](#)” on page 259.
- 9 After the storage server is created, you can do the following.
 - To configure a disk pool for this storage server, click **Next**.
See “[Configuring a Media Server Deduplication Pool](#)” on page 84.
 - To exit the wizard, click **Close**.

About NetBackup Media Server Deduplication Pools

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

Two types of deduplication pools exist, as follows:

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

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

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

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

Table 5-11 Deduplication pools for duplication or replication

Type	Requirements
Optimized duplication within the same NetBackup domain	<p>Optimized duplication in the same domain requires the following deduplication pools:</p> <ul style="list-style-type: none"> ■ At least one for the backup storage, which is the source for the duplication operations. The source deduplication pool is in one deduplication node. ■ Another to store the copies of the backup images, which is the target for the duplication operations. The target deduplication pool is in a different deduplication node.
Auto Image Replication to a different NetBackup domain	<p>Auto Image Replication deduplication pools can be either replication source or replication target. The replication properties denote the purpose of the deduplication pool. The deduplication pools inherit the replication properties from their volumes.</p> <p>See “About the replication topology for Auto Image Replication” on page 127.</p> <p>Auto Image Replication requires the following deduplication pools:</p> <ul style="list-style-type: none"> ■ At least one replication source deduplication pool in the originating domain. A replication source deduplication pool is one to which you send your backups. The backup images on the source deduplication pool are replicated to a deduplication pool in the remote domain or domains. ■ At least one replication target deduplication pool in a remote domain or domains. A replication target deduplication pool is the target for the duplication operations that run in the originating domain. <p>See “About NetBackup Auto Image Replication” on page 122.</p>

Configuring a Media Server Deduplication Pool

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

See [“Configuring a storage server for a Media Server Deduplication Pool”](#) on page 72.

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

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

- The type of disk pool (**PureDisk**).

- 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 91.

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

To configure a Media Server Deduplication Pool by using the wizard

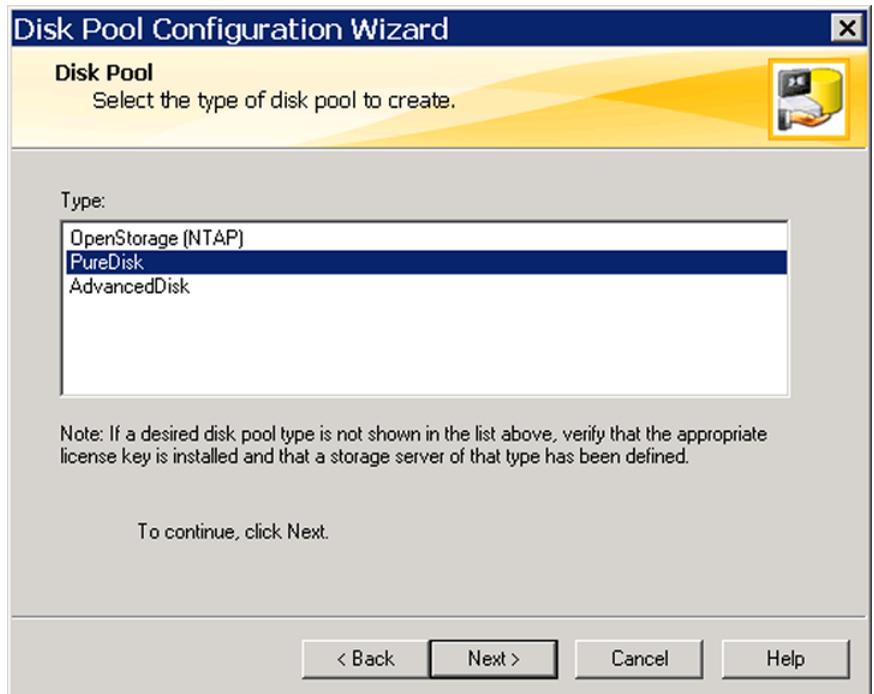
- 1 If the **Disk Pool Configuration Wizard** was launched from the **Storage Server Configuration Wizard**, go to step 3.

Otherwise, 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.

The following is an example of the wizard panel:

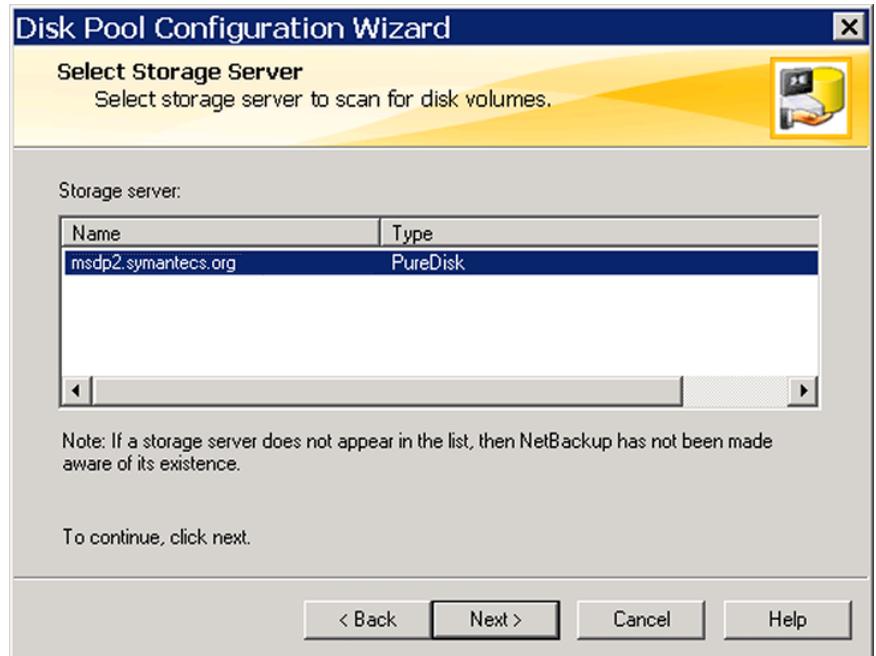


- 4 On the **Disk Pool** panel, select **PureDisk** in the **Type** windows.

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

Click **Next**. The **Select Storage Server** wizard panel appears.

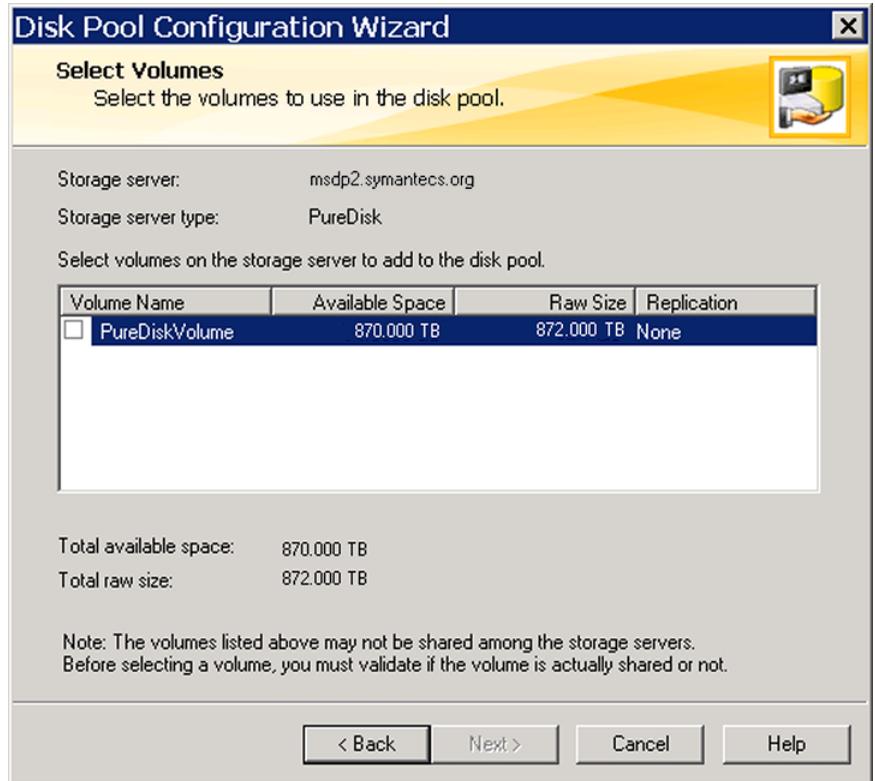
The following is an example of the wizard panel:



- 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.

Click **Next**. The **Select Volumes** wizard panel appears.

The following is an example of the wizard panel:

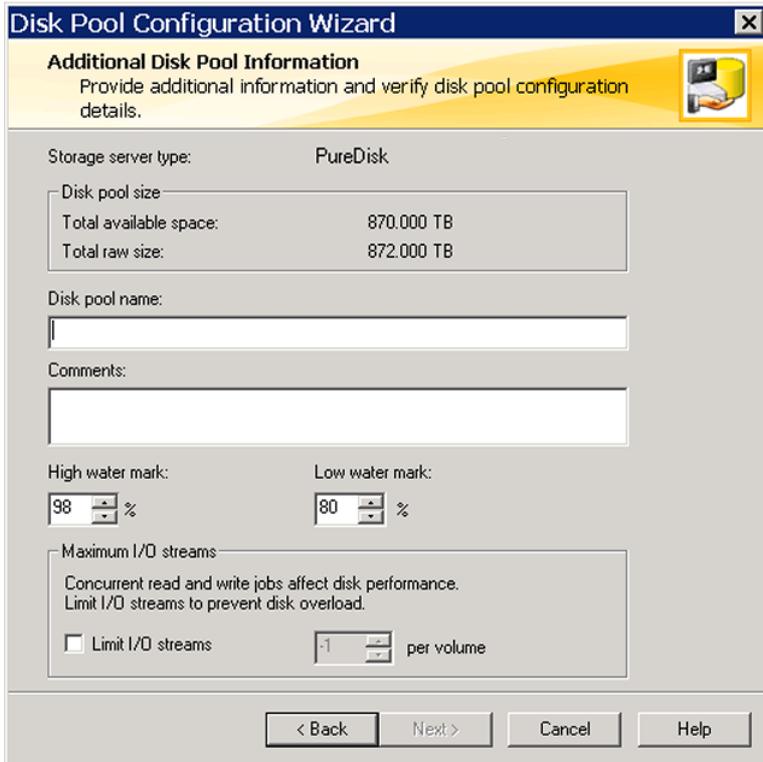


- 6 On the **Select Volumes** panel, select the volume for this disk pool.
 For a **PureDisk** type of disk pool, all disk storage is exposed as a single volume. The **PureDiskVolume** is a virtual name for the entire storage that is dedicated to the deduplicated backups.

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 a PureDisk storage pool.

Click **Next**. The **Additional Disk Pool Information** wizard panel appears. The following is an example of the wizard panel:



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

See “Media Server Deduplication Pool properties” on page 91.

Click **Next**. The **Summary** wizard panel appears.

The following is an example of the wizard panel:



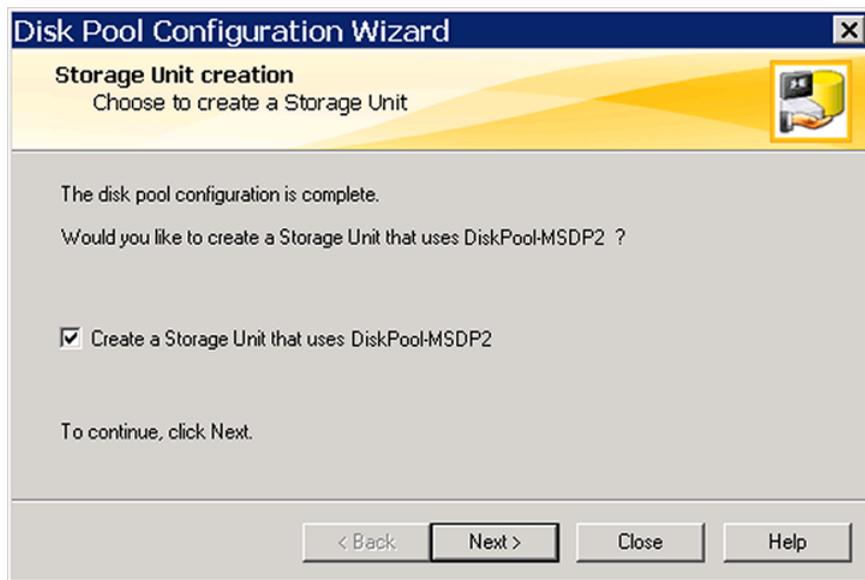
- 8 On the **Summary** panel, verify the selections. If OK, click **Next**.

- 9 After NetBackup creates the disk pool, a wizard panel describes the successful action.

To continue, click **Next**.

The **Storage Unit creation** wizard panel appears.

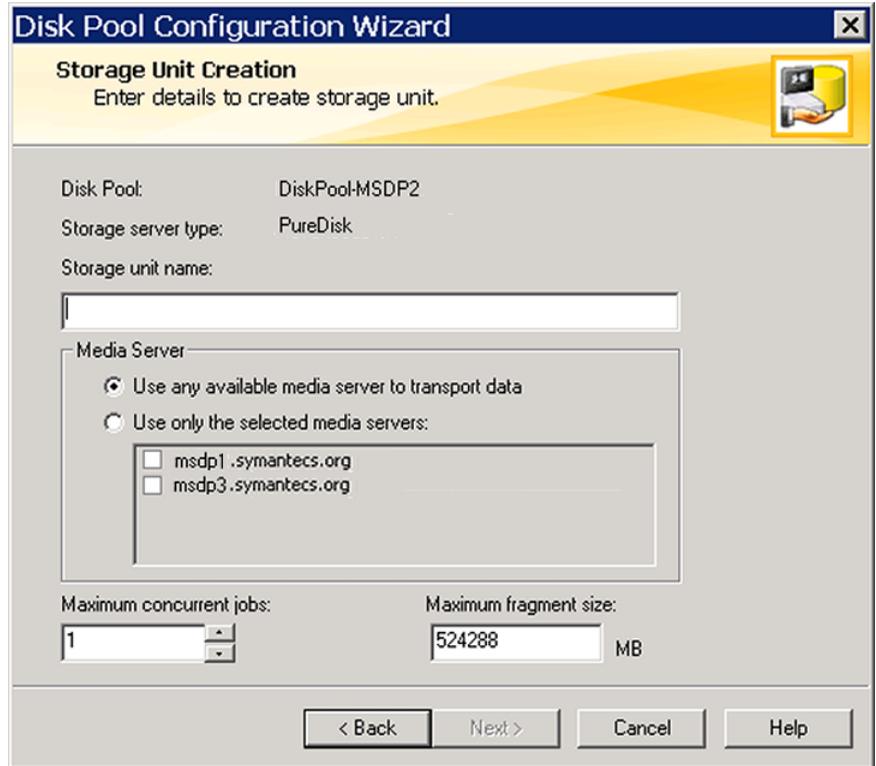
The following is an example of the wizard panel:



- To configure a storage unit for the disk pool, ensure that **Create a Storage Unit that uses *disk_pool_name*** is selected, then click **Next**. Otherwise, click **Close** to exit from the wizard.

If you click **Next**, a wizard panel appears in which you enter the details about the storage unit.

The following is an example of the wizard panel:



- Enter the appropriate information for the storage unit.
 See [“MSDP storage unit properties”](#) on page 95.
 Click **Next** to create the storage unit.
- After NetBackup configures the storage unit, the **Finished** panel appears. Click **Finish** to exit from the wizard.

Media Server Deduplication Pool properties

[Table 5-12](#) describes the disk pool properties.

Table 5-12 Media server deduplication pool properties

Property	Description
Name	The disk pool name.
Storage server	The storage server name. The storage server is the same as the NetBackup media server to which the storage is attached.
Disk volume	For a media server deduplication pool, all disk storage is exposed as a single volume. PureDiskVolume is a virtual name for the storage that is contained within the directories you specified for the storage path and the database path.
Refresh	Query the storage server for its replication capabilities. If the capabilities are different than the current NetBackup configuration, update the configuration.
Available space	The amount of space available in the disk pool.
Raw size	The total raw size of the storage in the disk pool.
Comment	A comment that is associated with the disk pool.
High water mark	The High water mark indicates that the volume is full. When the volume reaches the High water mark , NetBackup fails any backup jobs that are assigned to the storage unit. NetBackup also does not assign new jobs to a storage unit in which the deduplication pool is full. The High water mark includes the space that is committed to other jobs but not already used. The default is 98%. The default is 98%.
Low water mark	The Low water mark has no affect on the PureDiskVolume .

Table 5-12 Media server deduplication pool properties (*continued*)

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

Configuring an MSDP storage unit

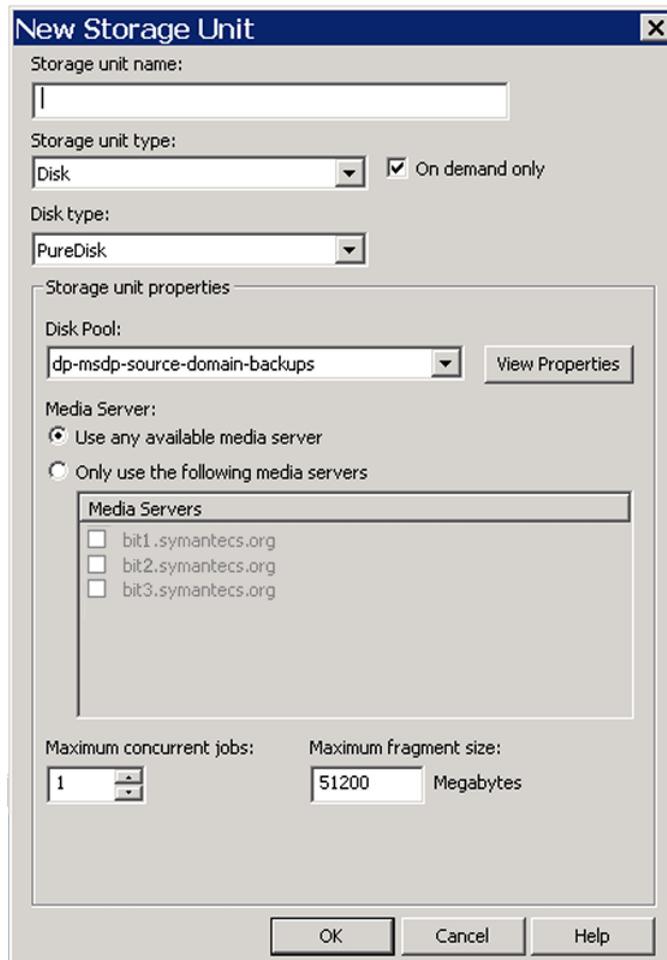
Create one or more storage units that reference the disk pool.

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

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

To configure a storage unit from the Actions menu

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



- 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 [“MSDP storage unit properties”](#) on page 95.

See “MSDP storage unit recommendations” on page 96.

MSDP storage unit properties

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

Table 5-13 MSDP storage unit properties

Property	Description
Storage unit name	A unique name for the new storage unit. The name can describe the type of storage. The storage unit name is the name used to specify a storage unit for policies and schedules. The storage unit name cannot be changed after creation.
Storage unit type	Select Disk as the storage unit type.
Disk type	Select PureDisk for the disk type for a Media Server Deduplication Pool , a PureDisk deduplication pool, or a PureDisk Deduplication Option storage pool.
Disk pool	Select the disk pool that contains the storage for this storage unit. All disk pools of the specified Disk type appear in the Disk pool list. If no disk pools are configured, no disk pools appear in the list.
Media server	The Media server setting specifies the NetBackup media servers that can deduplicate the data for this storage unit. Only the 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"> ■ To allow any server in the media server list to deduplicate data, select Use any available media server. ■ To use specific media servers to deduplicate the data, select Only use the following media servers. Then, select the media servers to allow. NetBackup selects the media server to use when the policy runs.
Maximum fragment size	For normal backups, NetBackup breaks each backup image into fragments so it does not exceed the maximum file size that the file system allows. You can enter a value from 20 MBs to 51200 MBs. For a FlashBackup policy, Symantec recommends that you use the default, maximum fragment size to ensure optimal deduplication performance. For more information, see the NetBackup Snapshot Client Administrator's Guide .

Table 5-13 MSDP storage unit properties (*continued*)

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

MSDP storage unit recommendations

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

Configure a client-to-server ratio

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

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

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

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

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

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

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

Throttle traffic to the media servers

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

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

Enabling MSDP encryption

Two procedures exist to enable encryption for MSDP, as follows:

- You can enable encryption on all hosts that deduplicate their own data without configuring them individually.
Use this procedure if you want all of your clients that deduplicate their own data to encrypt that data.
See [“To enable encryption on all hosts”](#) on page 98.
- You can enable encryption on individual hosts.

Use this procedure to enable compression or encryption on the storage server, on a load balancing server, or on a client that deduplicates its own data.

See [“To enable encryption on a single host”](#) on page 98.

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

To enable encryption on all hosts

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

To enable encryption on a single host

- 1 Use a text editor to open the `pd.conf` file on the host.
The `pd.conf` file resides in the following directories:
 - (UNIX) `/usr/openv/lib/ost-plugins/`
 - (Windows) `install_path\Veritas\NetBackup\bin\ost-plugins`See [“MSDP pd.conf file parameters”](#) on page 157.
- 2 For the line in the file that contains `ENCRYPTION`, remove the pound character (#) in column 1 from that line.
- 3 In that line, replace the 0 (zero) with a 1.

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

- 4 Ensure that the `LOCAL_SETTINGS` parameter is set to 1.
If `LOCAL_SETTINGS` is 0 (zero) and the `ENCRYPTION` setting on the storage server is 0, the client setting does not override the server setting. Consequently, the data is not encrypted on the client host.
- 5 Save and close the file.
- 6 If the host is the storage server or a load balancing server, restart the NetBackup Remote Manager and Monitor Service (`nbrmms`) on the host.

Configuring optimized synthetic backups for 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 35.

Table 5-14 To configure optimized synthetic backups for MSDP

Step	Task	Procedure
Step 1	If your storage server is at NetBackup release 7.0, set the OptimizedImage attribute on the storage server.	The OptimizedImage attribute is set by default beginning with the NetBackup 7.1 release. See “Setting MSDP storage server attributes” on page 195.
Step 2	If your storage server is at NetBackup release 7.0, set the OptimizedImage attribute on your existing deduplication pools. (Any deduplication pools that you create after you set the storage server attribute inherit the new functionality.)	The OptimizedImage attribute is set by default beginning with the NetBackup 7.1 release. See “Setting a Media Server Deduplication Pool attribute” on page 206.
Step 3	Configure a Standard or MS-Windows backup policy. Select the Synthetic backup attribute on the Schedule Attribute tab of the backup policy.	See “Creating a backup policy” on page 146.

See [“Setting MSDP storage server attributes”](#) on page 195.

See [“Creating a backup policy”](#) on page 146.

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 5-15 To configure optimized duplication of deduplicated data

Step	Action	Description
Step 1	Review the requirements and limitations	<p>See “About MSDP optimized duplication within the same domain” on page 103.</p> <p>See “About MSDP optimized duplication requirements” on page 104.</p> <p>See “About MSDP optimized duplication limitations” on page 104.</p>
Step 2	Configure the storage servers	<p>See “Configuring a storage server for a Media Server Deduplication Pool” on page 72.</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 “About the media servers for MSDP optimized duplication within the same domain” on page 106.</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 Disk Pool Configuration Wizard to configure them.</p> <p>See “Configuring a Media Server Deduplication Pool” on page 84.</p>
Step 4	Configure the storage unit for backups	<p>In the storage unit for your backups, do the following:</p> <ul style="list-style-type: none"> ■ For the Disk type, select PureDisk. ■ For the Disk pool, select one of the following: <ul style="list-style-type: none"> ■ If you back up to integrated NetBackup deduplication, select your Media Server Deduplication Pool. ■ If you back up to a PureDisk environment, select the PureDisk Deduplication Pool. <p>If you use a pull configuration, do not select the common media server in the backup storage unit. If you do, NetBackup uses it to deduplicate backup data. (That is, unless you want to use it for a load balancing server for the source deduplication node.)</p> <p>See “Configuring an MSDP storage unit” on page 93.</p>

Table 5-15 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> <ul style="list-style-type: none"> ■ For the Disk type, select PureDisk. ■ For the Disk pool, the destination can be a Media Server Deduplication Pool or a PureDisk Deduplication Pool. <p>Note: If the backup destination is a PureDisk Deduplication Pool, the duplication destination also must be a PureDisk Deduplication Pool.</p> <p>Also select Only use the following media servers. Then, select the media server or media servers that are common to both the source storage server and the destination storage server. If you select more than one, NetBackup assigns the duplication job to the least busy media server.</p> <p>If you select only a media server (or servers) that is not common, the optimized duplication job fails.</p> <p>See “Configuring an MSDP storage unit” on page 93.</p>
Step 6	Configure a separate network path for optimized duplication traffic	<p>Optionally, you can use a separate network for the optimized duplication traffic.</p> <p>See “Configuring a separate network path for MSDP optimized duplication” on page 115.</p>
Step 7	Configure optimized duplication behaviors	<p>Optionally, you can configure the optimized duplication behavior.</p> <p>See “Configuring NetBackup optimized duplication behavior” on page 117.</p> <p>See “About configuring MSDP optimized duplication and replication bandwidth” on page 119.</p>

Table 5-15 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"> ■ The first operation must be a Backup. For the Storage for the Backup operation, select the storage unit that is the target of your backups. That storage unit may use a Media Server Deduplication Pool or a PureDisk Deduplication Pool. These backups are the primary backup copies; they are the source images for the duplication operation. ■ For the second, child Operation, select Duplication. Then, select the storage unit for the destination deduplication pool. That pool may be a Media Server Deduplication Pool or a PureDisk Deduplication Pool. If the backup destination is a PureDisk Deduplication Pool, the duplication destination also must be a PureDisk Deduplication Pool. <p>See "About storage lifecycle policies" on page 137.</p> <p>See "Creating a storage lifecycle policy" on page 140.</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"> ■ If you use a storage lifecycle policy to manage the backup job and the duplication job: Select that storage lifecycle policy in the Policy storage field of the Policy Attributes tab. ■ If you do not use a storage lifecycle policy to manage the backup job and the duplication job: Select the storage unit for the Media Server Deduplication Pool that contains your normal backups. These backups are the primary backup copies. <p>See "About MSDP backup policy configuration" on page 145.</p> <p>See "Creating a backup policy" on page 146.</p>

Table 5-15 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"> ■ Configure a Vault profile. <ul style="list-style-type: none"> ■ On the Vault Profile dialog box Choose Backups tab, choose the backup images in the source Media Server Deduplication Pool. ■ On the Profile dialog box Duplication tab, select the destination storage unit in the Destination Storage unit field. ■ Configure a Vault policy to schedule the duplication jobs. A Vault policy is a NetBackup policy that is configured to run Vault jobs. <p>See the NetBackup Vault Administrator's Guide.</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 the source storage to the destination storage. The destination storage may be a Media Server Deduplication Pool or a PureDisk Deduplication Pool.</p> <p>See NetBackup Commands Reference Guide.</p>

About MSDP optimized duplication within the same domain

Optimized duplication within the same domain copies the backup images from a **Media Server Deduplication Pool** to deduplication storage in the same domain. The source and the destination storage must use the same NetBackup master server. The following table describes the supported source and targets.

Table 5-16 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

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.

Review the requirements and limitations.

See [“About MSDP optimized duplication requirements”](#) on page 104.

See [“About MSDP optimized duplication limitations”](#) on page 104.

About MSDP optimized duplication requirements

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

- If the source images reside on a NetBackup **Media Server Deduplication Pool**, the destination can be another **Media Server Deduplication Pool** or a **PureDisk Deduplication Pool**. (In NetBackup, a PureDisk storage pool is configured as a **PureDisk Deduplication Pool**.)
If the destination is a PureDisk storage pool, the PureDisk environment must be at release level 6.6 or later.
- If the source images reside on a PureDisk storage pool, the destination must be another PureDisk storage pool.
Both PureDisk environments must be at release level 6.6 or later.
- The source storage and the destination storage must have at least one media server in common.
See [“About the media servers for MSDP optimized duplication within the same domain”](#) on page 106.
- In the storage unit you use for the destination for the optimized duplication, you must select only the common media server or media servers.
If you select more than one, NetBackup assigns the duplication job to the least busy media server. If you select a media server or servers that are not common, the optimized duplication job fails.
For more information about media server load balancing, see the [NetBackup Administrator's Guide, Volume I](#).
- 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:

- You cannot use optimized duplication from a PureDisk storage pool (a **PureDisk Deduplication Pool**) to a **Media Server Deduplication Pool**.

- If an optimized duplication job fails after the configured number of retries, NetBackup does not run the job again.

By default, NetBackup retries an optimized duplication job three times. You can change the number of retries.

See [“Configuring NetBackup optimized duplication behavior”](#) on page 117.
- 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 storage pools, NetBackup reports that the deduplication rate is 100%. However, deduplication does not occur during optimized duplication. Therefore, ignore the deduplication rate statistics.
- 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 PureDisk storage, that PureDisk storage domain can only be the destination of MSDP duplication from one NetBackup domain. If you try to duplicate from MSDP storage in more than one NetBackup domain, optimized duplication does not function correctly.

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

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

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

About MSDP push duplication within the same domain

[Figure 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

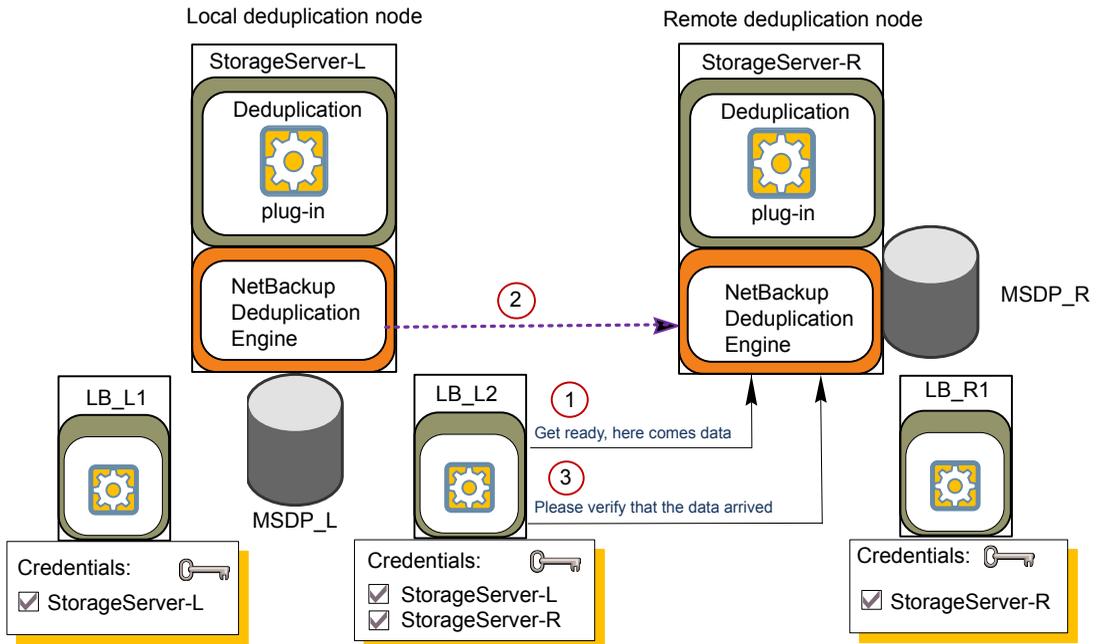


Figure 5-2 shows the settings for the storage unit for the normal backups for the local deduplication node. The disk pool is the **MSDP_L** in the local environment. Because all hosts in the local node are co-located, you can use any available media server for the backups.

Figure 5-2 Storage unit settings for backups to MSDP_L

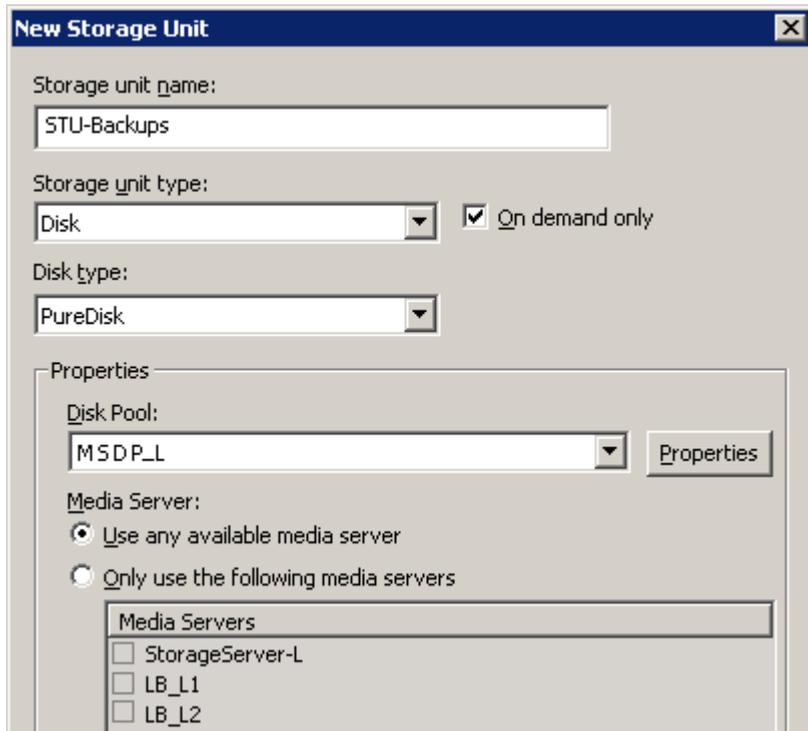
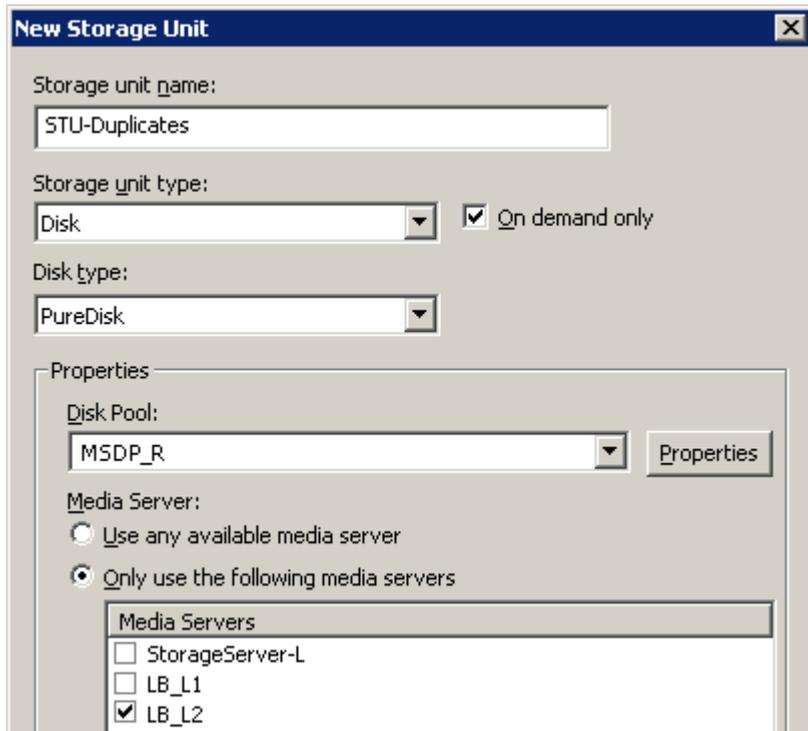


Figure 5-3 shows the storage unit settings for the optimized duplication. The destination is the **MSDP_R** in the remote environment. You must select the common server, so only load balancing server **LB_L2** is selected.

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 deduplication pool. In such a case, data travels across your WAN.

Figure 5-4 shows a push duplication from a **Media Server Deduplication Pool** to a PureDisk storage pool. The **Media Server Deduplication Pool** contains normal backups; the **PureDisk Deduplication Pool** is the destination for the optimized duplication copies. The StorageServer-A has credentials for both environments; it is the common media server.

Figure 5-4 Push duplication to a PureDisk storage pool

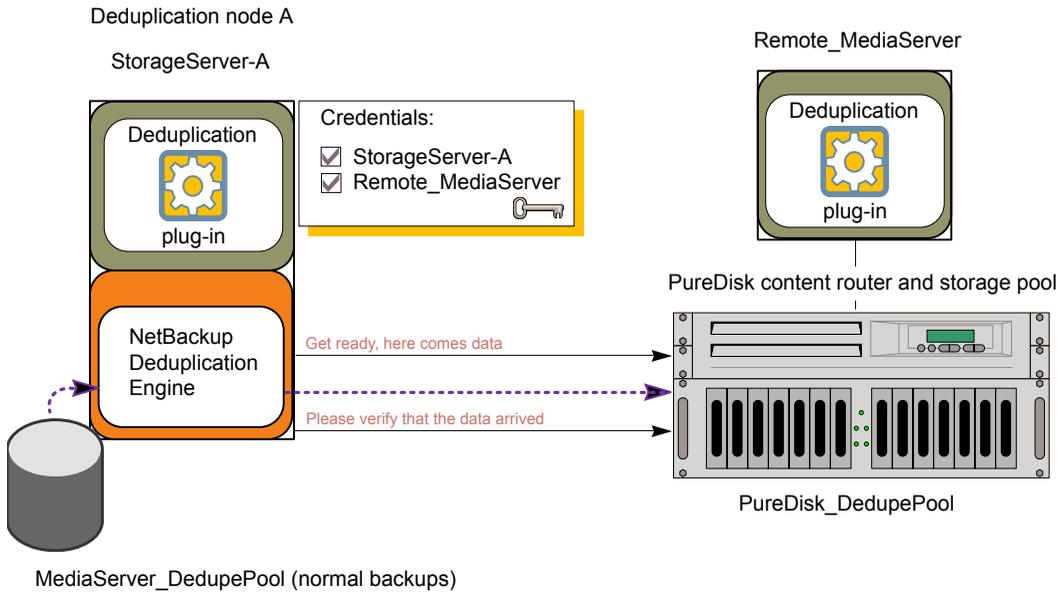


Figure 5-5 shows the storage unit settings for normal backups for the environment in Figure 5-4. The disk pool is the **MediaServer_DedupePool** in the local environment. For normal backups, you do not want a remote host deduplicating data, so only the local host is selected.

Figure 5-5 Storage unit settings for backups to MediaServer_DedupePool

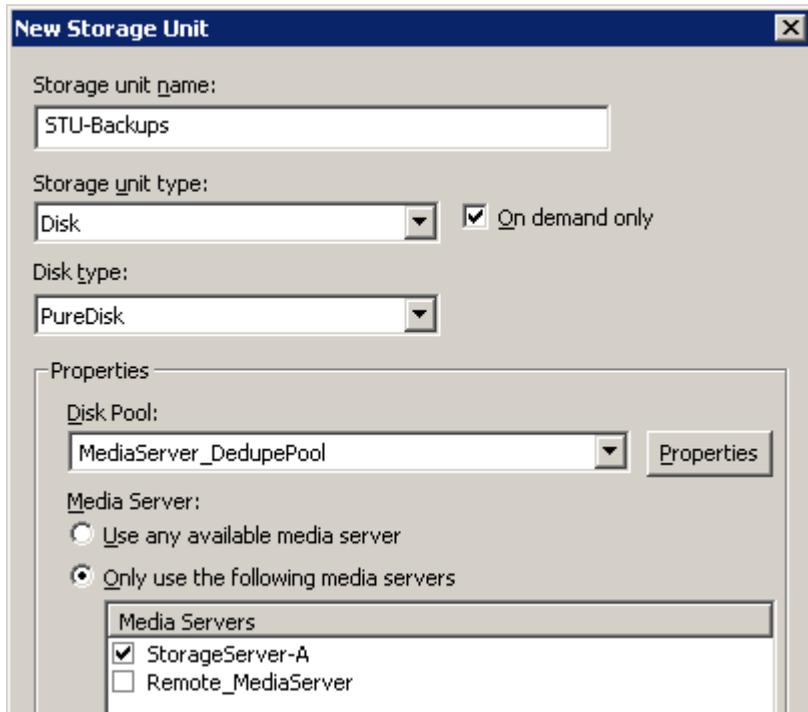


Figure 5-6 shows the storage unit settings for duplication for the environment in Figure 5-4. The disk pool is the **PureDisk_DedupePool** in the remote environment. You must select the common server, so only the local media server is selected. If this configuration were a pull configuration, the remote host would be selected in the storage unit.

Figure 5-6 Storage unit settings for duplication to PureDisk_DedupePool

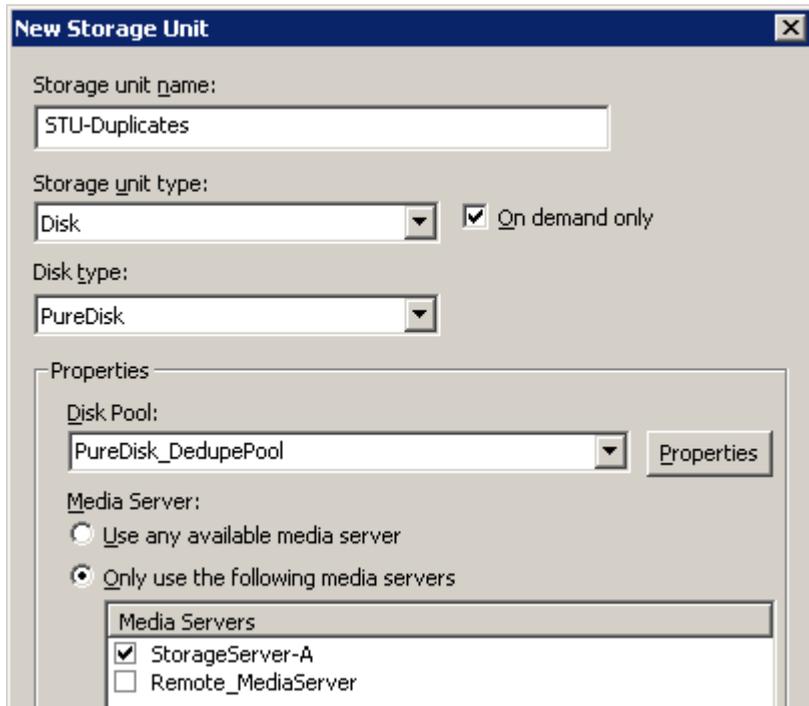
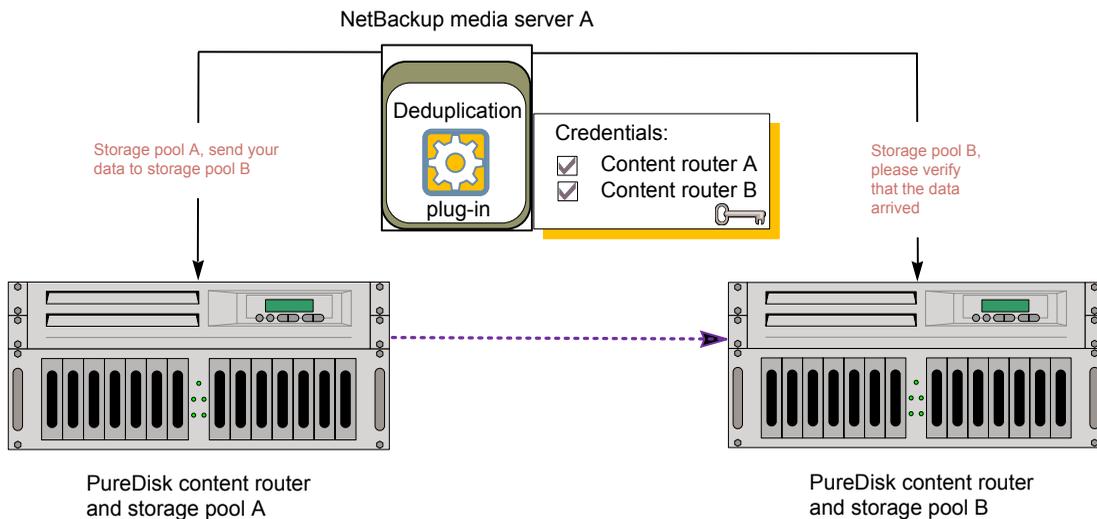


Figure 5-7 shows optimized duplication between two PureDisk storage pools. NetBackup media server A has credentials for both storage pools; it initiates, monitors, and verifies the optimized duplication. In the destination storage unit, the common server (media server A) is selected. This configuration is a push configuration. For a **PureDisk Deduplication Pool** (that is, a PureDisk storage pool), the PureDisk content router functions as the storage server.

Figure 5-7 Storage pool duplication



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

About MSDP pull duplication within the same domain

Figure 5-8 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 5-8 Pull duplication

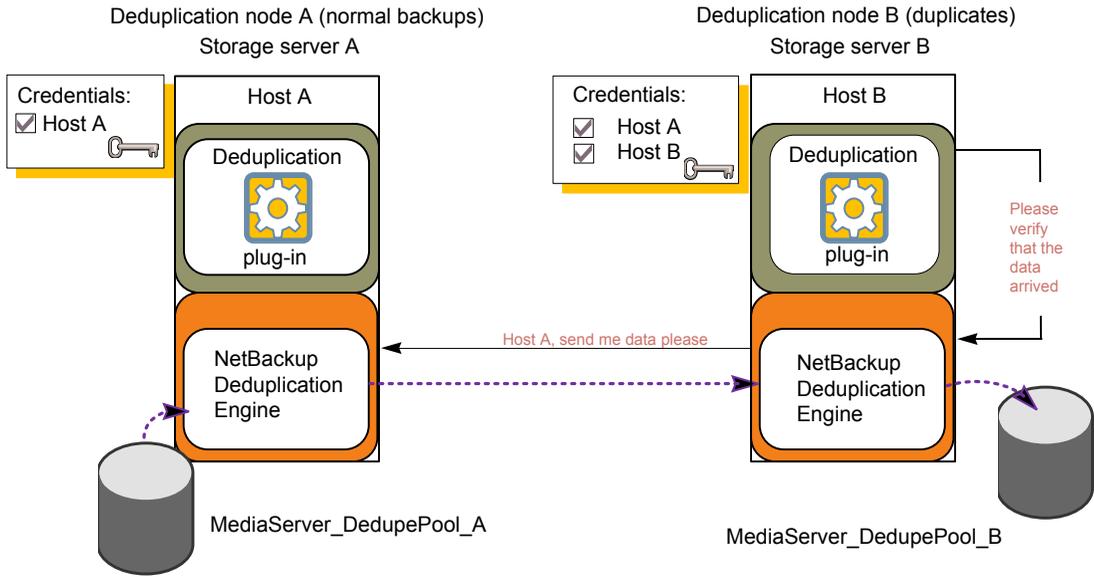
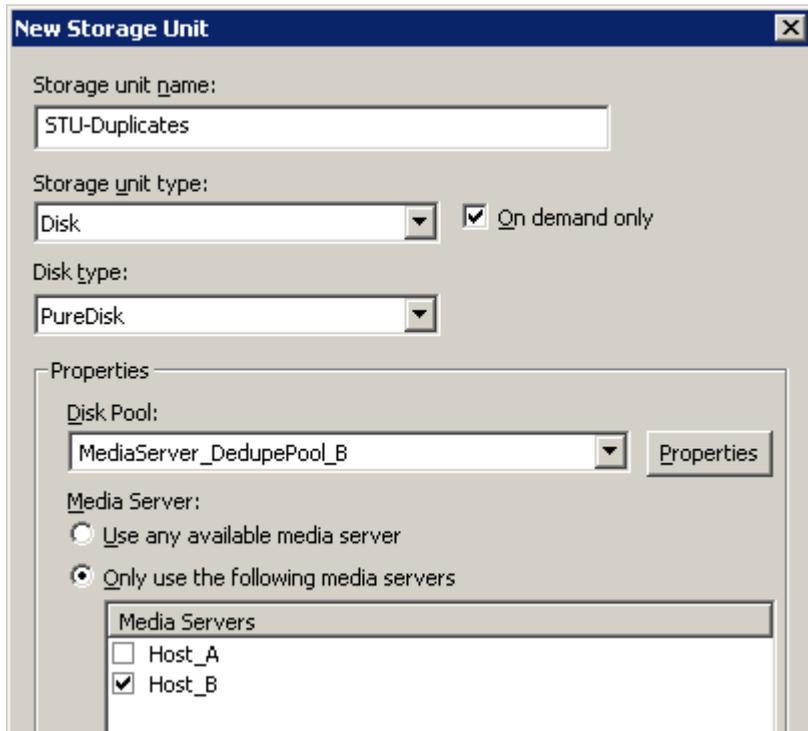


Figure 5-9 shows the storage unit settings for the duplication destination. They are similar to the push example except host B is selected. Host B is the common server, so it must be selected in the storage unit.

Figure 5-9 Pull duplication storage unit settings



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

Configuring a separate network path for MSDP optimized duplication

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

The following are the requirements:

- A separate, dedicated network interface card in both the source and the destination storage servers.
- The separate network is operational and using the dedicated network interface cards on the source and the destination storage servers.

To configure a separate network path for MSDP optimized duplication

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

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

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

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

```
hosts:                files dns
```

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

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

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

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

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

```
hosts:                files dns
```

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

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

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

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

Configuring NetBackup optimized duplication behavior

You can configure the following optimized duplication behaviors for NetBackup:

Table 5-17 Optimized duplication behavior

Behavior	Description
Optimized duplication failover	<p>By default, if an optimized duplication job fails, NetBackup does not run the job again.</p> <p>You can configure NetBackup to use normal duplication if an optimized duplication fails.</p> <p>For example, NetBackup does not support optimized duplication from a PureDisk Storage Pool Authority source to a Media Server Deduplication Pool. Therefore, to migrate data from a PureDisk SPA into a Media Server Deduplication Pool, you must change the default NetBackup failover behavior for optimized duplication.</p> <p>See "To configure optimized duplication failover" on page 118.</p>
Number of optimized duplication attempts	<p>By default, NetBackup tries an optimized duplication job three times before it fails the job.</p> <p>You can change the number of times NetBackup retries an optimized duplication job before it fails the jobs.</p> <p>See "To configure the number of duplication attempts" on page 118.</p>
Storage lifecycle policy retry wait period	<p>If a storage lifecycle policy optimized duplication job fails, NetBackup waits two hours and then retries the job. By default, NetBackup tries a job three times before the job fails.</p> <p>You can change the number of hours for the wait period.</p> <p>See "To configure the storage lifecycle policy wait period" on page 118.</p>

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

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 118.

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

To configure the number of duplication attempts

- ◆ On the master server, create a file named `OPT_DUP_BUSY_RETRY_LIMIT` that contains an integer that specifies the number of times to retry the job before NetBackup fails the job.

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

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

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.symantecs.org`

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/openv/netbackup/bin/nbsetconfig`

UNIX NetBackup servers: `/usr/openv/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 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 156.

See [“Editing the MSDP pd.conf file”](#) on page 156.

See [“MSDP pd.conf file parameters”](#) on page 157.

See [“Configuring MSDP optimized duplication within the same NetBackup domain”](#) on page 99.

See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 120.

Configuring MSDP replication to a different NetBackup domain

[Table 5-18](#) describes the tasks that are required to replicate backup images from one **Media Server Deduplication Pool** to another in a different NetBackup domain.

Table 5-18 NetBackup MSDP replication configuration tasks

Step	Task	Procedure
Step 1	Learn about NetBackup Auto Image Replication	See “About NetBackup Auto Image Replication” on page 122.
Step 2	Understand the rules for replication	See “About MSDP replication to a different domain” on page 122.
Step 3	Determine if you need to configure a trust relationship with the target NetBackup domain	A trust relationship is optional. See “About trusted master servers for Auto Image Replication” on page 130.
Step 4	Add the remote storage server as a replication target	See “Configuring a target for MSDP replication to a remote domain” on page 132. See “Viewing the replication topology for Auto Image Replication” on page 128.
Step 5	Configure a storage lifecycle policy	<p>The following are the options when you configure the SLP operations:</p> <ul style="list-style-type: none"> ■ If you configured a trust relationship with the target domains, you can specify one of the following options: <ul style="list-style-type: none"> ■ All replication target storage servers (across different NetBackup domains) NetBackup automatically creates an import SLP in the target domain when the replication job runs. ■ A specific Master Server. If you choose this option, you then select Target master server and Target import SLP. You must create an import SLP in the target domain before you configure an SLP in the source domain. ■ If you did <i>not</i> configure a trust relationship with the target domains, All replication target storage servers (across different NetBackup domains) 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. <p>See “About storage lifecycle policies” on page 137.</p> <p>See “About the storage lifecycle policies required for Auto Image Replication” on page 138.</p> <p>See “Creating a storage lifecycle policy” on page 140.</p>

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 5-19](#) describes the MSDP replication source and targets that NetBackup supports.

Table 5-19 NetBackup media server deduplication replication targets

Source storage	Target storage
Media Server Deduplication Pool on a NetBackup media server	A Media Server Deduplication Pool on a NetBackup media server.
Media Server Deduplication Pool on a NetBackup media server	A PureDisk Deduplication Pool on a NetBackup media server. A PureDisk Deduplication Pool is a PureDisk Storage Pool exposed to NetBackup by PureDisk.
Media Server Deduplication Pool on a NetBackup 5200 series appliance	A Media Server Deduplication Pool 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. 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.

If the destination is PureDisk storage, that PureDisk storage domain can only be the destination of MSDP duplication from one NetBackup domain. If you try to duplicate from MSDP storage in more than one NetBackup domain, optimized duplication does not function correctly.

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 124.
- 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 5-20 NetBackup Auto Image Replication storage types

Storage type	Link to more information
Media Server Deduplication Pool	See “About MSDP replication to a different domain” on page 122.

Notes about Auto Image Replication

- Although Auto Image Replication is a disaster recovery solution, the administrator cannot directly restore to clients in the primary (or originating) domain from the target master domain.
- 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.
- Replication between the source domain and the target domain must be between supported versions of NetBackup. For example, replicating from NetBackup 7.6 in the source domain to NetBackup 7.0 in the target domain can lead to data loss.
- 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 5-21](#) 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 138.

Table 5-21 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 Policy storage selection. The SLP must include at least one Replication 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 Catalog 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

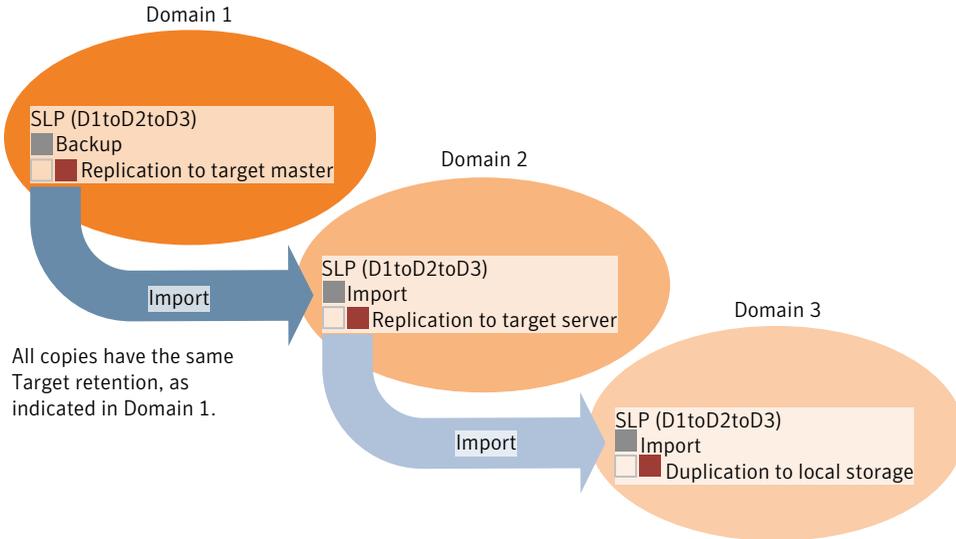
Replications can be cascaded from the originating domain to multiple domains. Storage lifecycle policies are set up in each domain to anticipate the originating image, import it and then replicate it to the next target master.

[Figure 5-10](#) 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 5-10 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 5-22 shows the difference in the import operation setup.

Table 5-22 Import operation difference in an SLP configured to replicate the imported copy

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

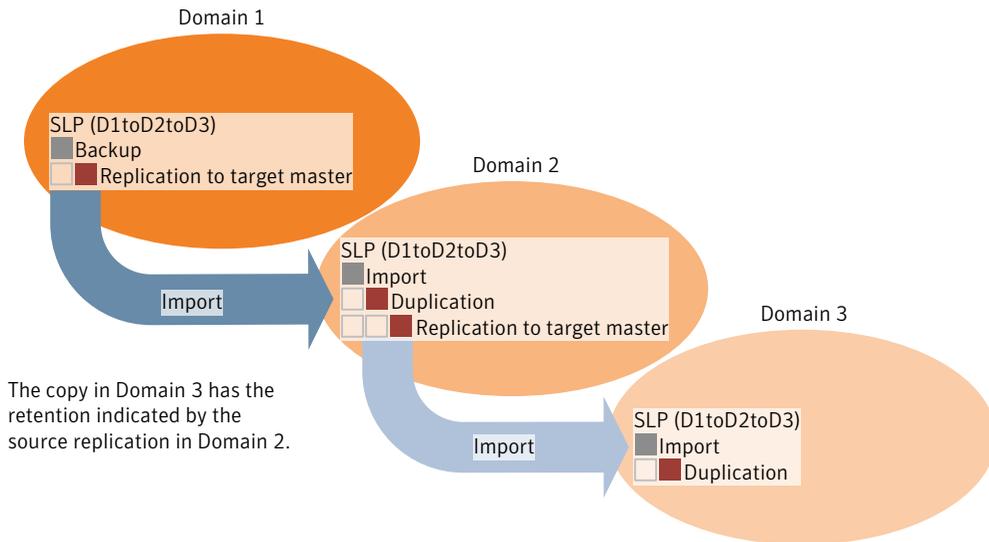
The target retention is embedded in the source image.

Because the imported copy is replicated to a target master server domain, NetBackup uses the target retention period rather than .

In the cascading model that is represented in [Figure 5-10](#), 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 5-11 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 132.

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 130.

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:

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

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 132.

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 207.

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 128.

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 127.

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:

`-storage_server host_name` The name of the target storage server.

`-stype PureDisk` For a **Media Server Deduplication Pool** target or a **PureDisk Deduplication Pool** target, use `PureDisk`.

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 129.

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.

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 137.

See [“Creating a storage lifecycle policy”](#) on page 140.

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 131.

See [“About NetBackup Auto Image Replication”](#) on page 122.

Enabling NetBackup clustered master server inter-node authentication

Before you can configure trust relationships for a NetBackup master server that is clustered, you must enable inter-node communication on all of the nodes in the cluster. Do so before you add the trusted master server. This requirement applies regardless of whether the clustered mastered server is the source of the replication operation or the target.

To enable clustered master server inter-node authentication

- ◆ On each node of the clustered NetBackup master server, run the following NetBackup command:
 - **Windows:** `install_path\Veritas\NetBackup\bin\admincmd\bpnbaz -setupat`
 - **UNIX:** `/usr/opensv/netbackup/bin/admincmd/bpnbaz -setupat`

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.symantecs.org'
Setting up security on target host: bitl.remote.symantecs.org
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 122.

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 120.

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 130.

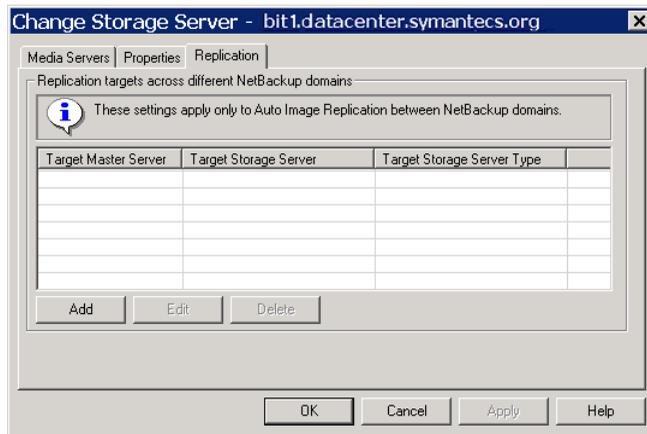
See [“Enabling NetBackup clustered master server inter-node authentication”](#) on page 131.

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

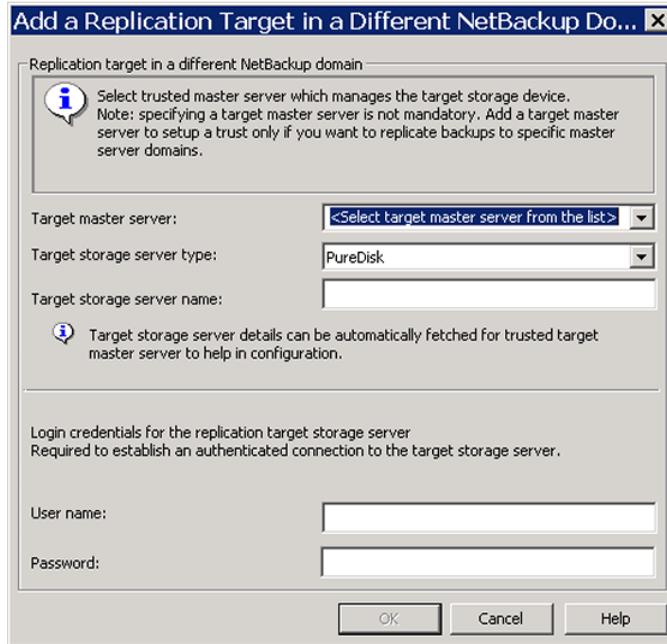
- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Credentials > Storage Server**.
- 2 Select the MSDP storage server.
- 3 On the **Edit** menu, select **Change**.
- 4 In the **Change Storage Server** dialog box, select the **Replication** tab.

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.



- 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.
- 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 130.

See [“Target options for MSDP replication”](#) on page 136.

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 136.

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 136.

- 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 207.

Target options for MSDP replication

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

Table 5-23 MSDP target replication options

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

Table 5-23 MSDP target replication options (*continued*)

Option	Description
Target storage server type	The target storage server type must be PureDisk . If a trusted master server is configured, PureDisk 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 Target storage server name field.
Target storage server name	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. The drop-down list shows all the storage servers that match the Target storage server type .
User name	When you configure a replication target, NetBackup populates the User name field with user account of the target storage server, as follows: <ul style="list-style-type: none">■ For an MSDP target, the NetBackup Deduplication Engine user name.■ For a PureDisk storage target, the PureDisk Storage Pool Authority user name.
Password	Enter the password for the NetBackup Deduplication Engine.

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 5-24 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"> ■ The first operation must be a Backup operation to a Media Server Deduplication Pool. Indicate the exact storage unit from the drop-down list. Do not select Any Available. Note: The target domain must contain the same type of storage to import the image. ■ At least one operation must be a Replication operation to a Media Server Deduplication Pool or a PureDisk Deduplication Pool in another NetBackup domain. You can configure multiple Replication operations in an Auto Image Replication SLP. The Replication operation settings determine whether the backup is replicated to all replication targets in all master server domains or only to specific replication targets. See “About trusted master servers for Auto Image Replication” on page 130. ■ The SLP must be of the same data classification as the Import SLP in Domain 2. Note: 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 Any. Use a different data classification in the source domain or the Import job fails.

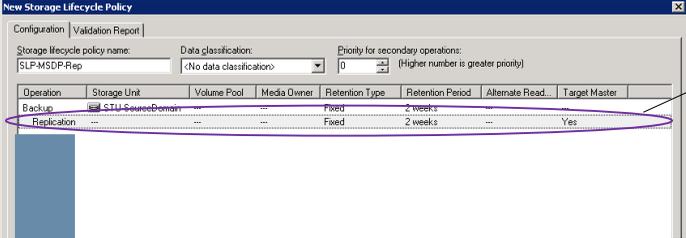
Table 5-24 SLP requirements for Auto Image Replication (*continued*)

Domain	Storage lifecycle policy requirements
Domain 2 (Target domain)	<p>If replicating to all targets in all domains, in each domain NetBackup automatically creates an Import SLP that meets all the necessary criteria.</p> <p>Note: If replicating to specific targets, you must create the Import SLP before creating the Auto Image Replication SLP in the originating domain.</p> <p>The Import SLP must meet the following criteria:</p> <ul style="list-style-type: none"> ■ The first operation in the SLP must be an Import operation. NetBackup must support the Destination storage as a target for replication from the source storage. Indicate the exact storage unit from the drop-down list. Do not select Any Available. ■ The SLP must contain at least one operation that has the Target retention specified. ■ 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. <p>See the following topic for more information about Replication operation configuration:</p>

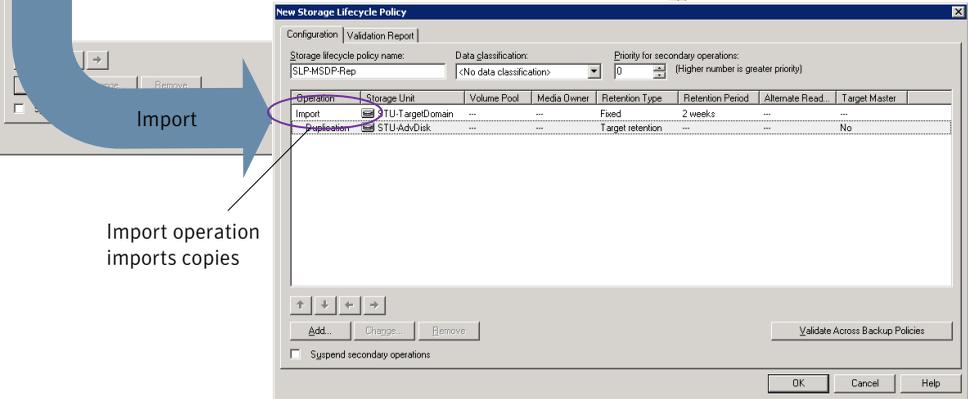
[Figure 5-12](#) shows how the SLP in the target domain is set up to replicate the images from the originating master server domain.

Figure 5-12 Storage lifecycle policy pair required for Auto Image Replication

SLP on master server in the source domain



Replication operation indicates a target master



Import operation imports copies

SLP that imports the copies to the target domain

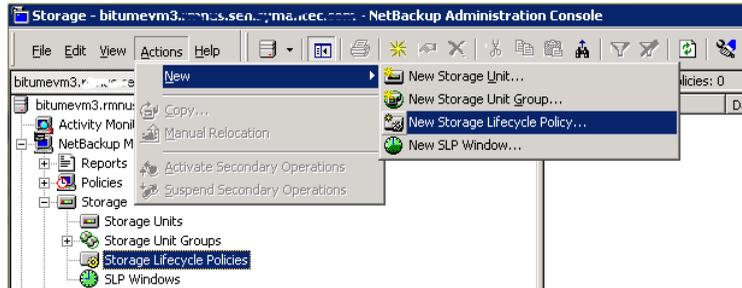
Note: Restart `nbstserv` after making changes to the underlying storage for any operation in an SLP.

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**.

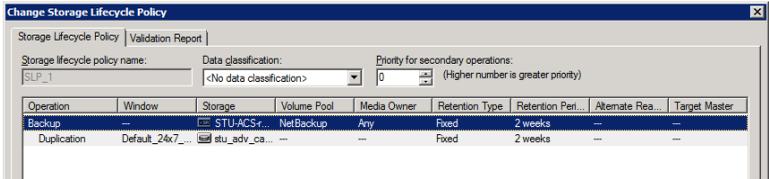
See [“NetBackup naming conventions”](#) on page 21.

- 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.

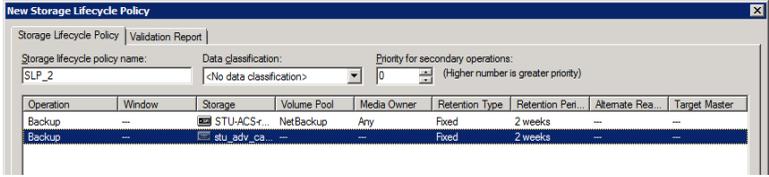
If this is the first operation added to the SLP, click **Add**.

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 Blue Replication](#)” on page 122.

- 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 146.

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 5-13 Storage Lifecycle Policy tab

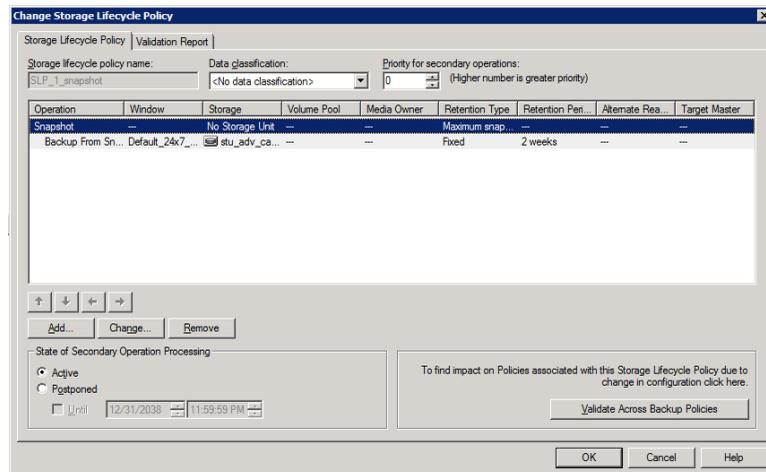


Table 5-25 Storage Lifecycle Policy tab

Setting	Description
Storage lifecycle policy name	The Storage lifecycle policy name describes the SLP. The name cannot be modified after the SLP is created.
Data classification	<p>The Data classification defines the level 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 Any classification, which is unique to SLPs.</p> <p>The Any 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 “About the storage lifecycle policies required for Auto Image Replication” on page 138.</p> <p>The Data classification is an optional setting.</p> <p>One data classification can be assigned to each SLP and applies to all operations in the SLP.</p> <p>If a data classification is selected (other than Any), the SLP stores only those images from the policies that are set up for that data classification. If no data classification is indicated, the SLP accepts images of any classification or no classification.</p> <p>The Data classification setting allows the NetBackup administrator to classify data based on relative importance. A classification represents a set of backup requirements. When data must meet different backup requirements, consider assigning different classifications.</p> <p>For example, email backup data can be assigned to the silver data classification and financial data backup may be assigned to the platinum classification.</p> <p>A backup policy associates backup data with a data classification. Policy data can be stored only in an SLP with the same data classification.</p> <p>Once data is backed up in an SLP, the data is managed according to the SLP configuration. The SLP defines what happens to the data from the initial backup until the last copy of the image has expired.</p>
Priority for secondary operations	<p>The Priority for secondary operations 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 Backup and Snapshot operations. Range: 0 (default) to 99999 (highest priority).</p> <p>For example, you may want to set the Priority for secondary operations 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 Attributes tab.</p>

Table 5-25 Storage Lifecycle Policy tab (*continued*)

Setting	Description
Operations	<p>Use the Add, Change, and Remove 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	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.
Active and Postponed	<p>The Active and Postponed options appear under State of Secondary Operation Processing and refer to the processing of all duplication operations in the SLP.</p> <p>Note: The Active and Postponed options apply to duplication operations that create tar-formatted images. For example, those created with <code>bpduplicate</code>. The Active and Postponed 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.</p> <ul style="list-style-type: none"> ■ Enable Active to let secondary operations continue as soon as possible. When changed from Postponed to Active, NetBackup continues to process the images, picking up where it left off when secondary operations were made inactive. ■ Enable Postponed to postpone the secondary operations for the entire SLP. Postponed 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. <p>All secondary operations in the SLP are inactive indefinitely unless the administrator selects Active or until the Until option is selected and an activation date is indicated.</p>
Validate Across Backup Policies 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 Validation Report tab.</p> <p>This button performs the same validation as the <code>-conflict</code> option performs when used with the <code>nbstl</code> command.</p>

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 21.
- 4 If necessary, clear the **Use Policy Configuration Wizard** checkbox.
- 5 Click **OK**.
- 6 Configure the attributes, the schedules, the clients, and the backup selections for the new policy.

Enabling MSDP client-side deduplication

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

To specify the clients that deduplicate backups

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the details pane, select the master server.
- 3 On the **Actions** menu, select **Properties**.
- 4 Select the **Client Attributes** properties.
- 5 On the **General** tab of the **Client Attributes** properties, add the clients that use client direct 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](#).

Configuring MSDP client-direct restores

Beginning with the NetBackup 7.6 release, NetBackup can restore directly to a client that deduplicates its own data. Previously, NetBackup used the media server components for the restore jobs. The media server components reside on the storage server or a load balancing server. Now, the restore jobs bypass the media server components. Client deduplication is also known as *client-side deduplication* or *client-direct*.

See “[About NetBackup MSDP Client Deduplication](#)” on page 28.

To enable restores directly to a client that deduplicates its own data

- 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. How you set the option depends on the operating system, as follows:
 - UNIX: Use a text editor to add it to the `/usr/opensv/netbackup/bp.conf` file.
 - Windows: Add it to the Windows registry.
See “[Setting NetBackup configuration options by using the command line](#)” on page 118.
- 2 On the master server, run the following command to configure NetBackup to use client-direct restores for the client:

```
UNIX: /usr/opensv/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
```

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**.

Resilient Network properties

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

The **Resilient Network** properties let you configure NetBackup to use resilient network connections. A resilient connection allows backup and restore traffic between a client and 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 (`nbrntd`) creates. Examples of the connections that are not supported are:

- Clients that back up their own data (deduplication clients and SAN clients)
- Granular Recovery for Exchange or SharePoint
- NetBackup `nbfsd` 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 5-14 Master server Resilient Network host properties

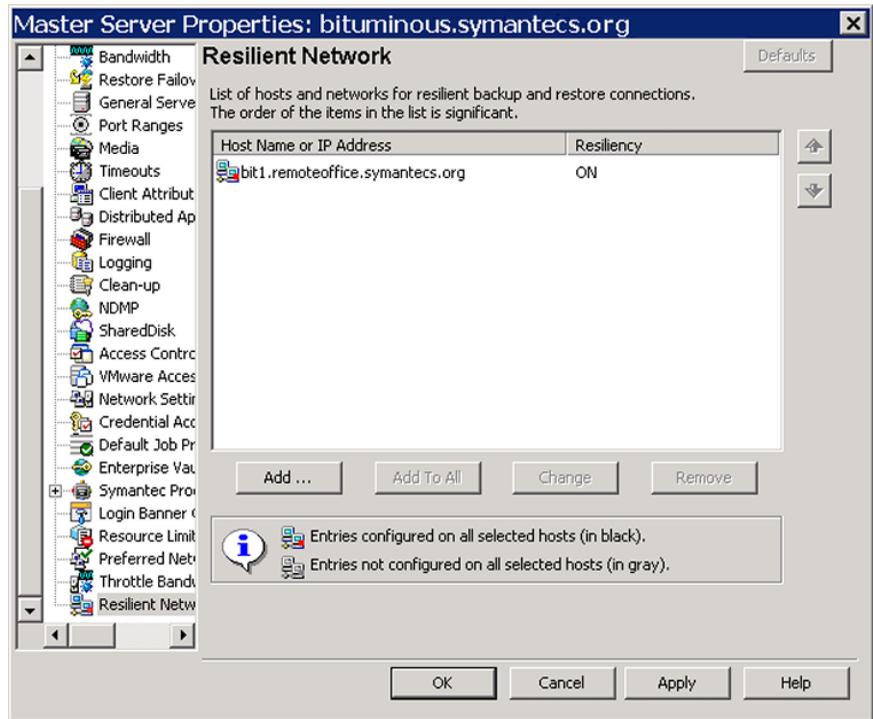


Table 5-26 describes the **Resilient Network** properties.

Table 5-26 Resilient Network dialog box properties

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

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

The resilient status of each client also appears as follows:

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

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

The NetBackup resilient connections use the SOCKS protocol version 5.

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

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

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 149.

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 finished 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 range as follows

- 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 the item up or down.

The order of the items in the list is significant.

See [“Resilient Network properties”](#) on page 149.

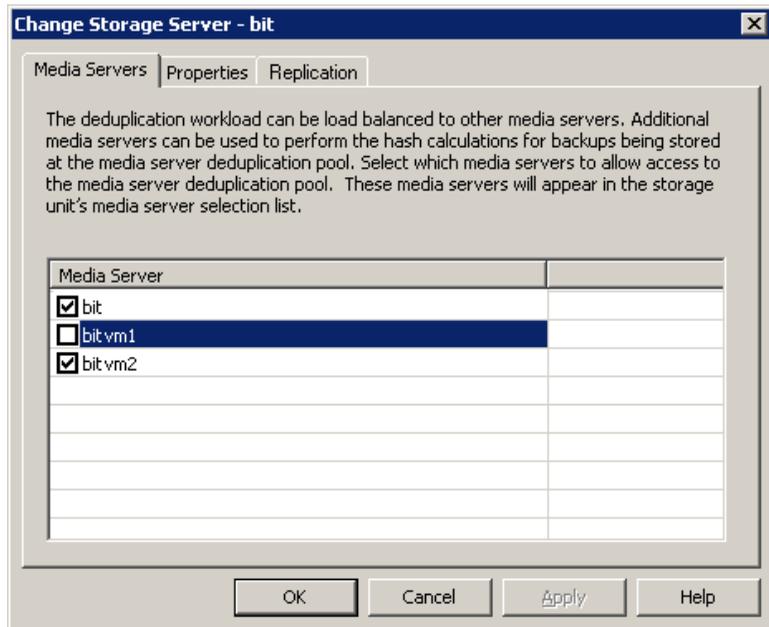
Adding an MSDP load balancing server

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

See [“About MSDP storage servers”](#) on page 25.

To add a load balancing server

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



- 4 In the **Change Storage Server** dialog box, select the **Media Servers** tab
- 5 Select the media server or servers that you want to use as a load balancing server. It must be a supported host.

The media servers that are checked are configured as load balancing servers.

- 6 Click **OK**.
- 7 For all storage units in which **Only use the following media servers** is configured, ensure that the new load balancing server is selected.

About the 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 157.

See [“Editing the MSDP `pd.conf` file”](#) on page 156.

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 156.

See [“MSDP `pd.conf` file parameters”](#) on page 157.

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 5-27](#) 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 5-27 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"> ■ Classless Inter-Domain Routing (CIDR) format. For example, the following notation specifies 192.168.10.0 and 192.168.10.1 for traffic: <code>BACKUPRESTORERANGE = 192.168.10.1/31</code> ■ Comma-separated list of IP addresses. For example, the following notation specifies 192.168.10.1 and 192.168.10.2 for traffic: <code>BACKUPRESTORERANGE = 192.168.10.1, 192.168.10.2</code> <p>Default value: <code>BACKUPRESTORERANGE=</code> (no default value)</p> <p>Possible values: Classless Inter-Domain Routing format notation or comma-separated list of IP addresses</p>
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.</p> <p>By default, files are not compressed.</p> <p>Default value: <code>COMPRESSION = 1</code></p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About MSDP compression” on page 33.</p>

Table 5-27 pd.conf file parameters (continued)

Parameter	Description
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"> ■ If disabled (set to 0), a request for the latest storage capacity information occurs whenever NetBackup requests it. ■ If you specify a value, a request occurs only after the specified number of seconds since the last request. Otherwise, a cached value from the previous request is used. ■ Enabling this setting may reduce the queries to the storage server. The drawback is the capacity information reported by NetBackup becomes stale. Therefore, if storage capacity is close to full, Symantec recommends that you do not enable this option. ■ On high load systems, the load may delay the capacity information reporting. If so, NetBackup may mark the storage unit as down. <p>Default value: CR_STATS_TIMER = 0</p> <p>Possible values: 0 or greater, in seconds</p>
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>Note: This parameter does not apply for NDMP backups from a NetApp appliance.</p> <p>Default value:</p> <ul style="list-style-type: none"> ■ UNIX: DEBUGLOG = /var/log/puredisk/pdplugin.log ■ Windows: DEBUGLOG = C:\pdplugin.log <p>Possible values: Any path</p>

Table 5-27 pd.conf file parameters (continued)

Parameter	Description
DONT_SEGMENT_TYPES	<p>A comma-separated list of file name extensions of files not to be deduplicated. Files in the backup stream that have the specified extensions are given a single segment if smaller than 16 MB. Larger files are deduplicated using the maximum 16-MB segment size.</p> <p>Example: <code>DONT_SEGMENT_TYPES = mp3,avi</code></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: <code>DONT_SEGMENT_TYPES = (no default value)</code></p> <p>Possible values: comma-separated file extensions</p>
ENCRYPTION	<p>Specifies whether to encrypt the data. By default, files are not encrypted..</p> <p>If you set this parameter to 1 on all hosts, the data is encrypted during transfer and on the storage.</p> <p>Default value: <code>ENCRYPTION = 0</code></p> <p>Possible values: 0 (no encryption) or 1 (encryption)</p> <p>See “About MSDP encryption” on page 34.</p>
FIBRECHANNEL	<p>Enable Fibre Channel for backup and restore traffic to and from a NetBackup series appliance.</p> <p>See “About Fibre Channel to a NetBackup 5020 appliance” on page 279.</p> <p>Default value: <code>FIBRECHANNEL = 0</code></p> <p>Possible values: 0 (off) or 1 (on)</p>

Table 5-27 pd.conf file parameters (*continued*)

Parameter	Description
<p><code>FILE_KEEP_ALIVE_INTERVAL</code></p>	<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"> ■ 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 “Resilient Network properties” on page 149. ■ Commented out (the default) and Resilient Network connections are <i>not</i> enabled: The keep-alive interval is 1440 seconds (80% of 30 minutes). ■ 0 or less: Disabled; no keepalives are sent. ■ 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). <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> <pre>Using keepalive interval of xxxx seconds</pre> <p>For more information about the deduplication plug-in log file, see <code>DEBUGLOG</code> and <code>LOGLEVEL</code> in this table.</p>

Table 5-27 pd.conf file parameters (continued)

Parameter	Description
FP_CACHE_CLIENT_POLICY	<p>Note: Symantec recommends that you use this setting on the individual clients that back up their own data (client-side deduplication). If you use it on a storage server or load balancing server, it affects all backup jobs.</p> <p>Specifies the client, backup policy, and date from which to obtain the fingerprint cache for the first backup of a client.</p> <p>By default, the fingerprints from the previous backup are loaded. This parameter lets you load the fingerprint cache from another, similar backup. It can reduce the amount of time that is required for the first backup of a client. This parameter especially useful for remote office backups to a central datacenter in which data travels long distances over a WAN.</p> <p>Specify the setting in the following format:</p> <p><i>clienthostmachine,backuppolicy,date</i></p> <p>The date is the last date in mm/dd/yyyy format to use the fingerprint cache from the client you specify.</p> <p>Default value: FP_CACHE_CLIENT_POLICY = (no default value)</p> <p>See “Seeding the MSDP fingerprint cache from the client” on page 70.</p>
FP_CACHE_INCREMENTAL	<p>Specifies whether to use fingerprint caching for incremental backups.</p> <p>Because incremental backups only back up what has changed since the last backup, cache loading has little affect on backup performance for incremental backups.</p> <p>Default value: FP_CACHE_INCREMENTAL = 0</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>
FP_CACHE_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 5-27 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>Note: 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>Note: 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"> ■ The container has not been rebased within the last three months. ■ For that backup, the data segments in the container consume less space than the <code>FP_CACHE_PERIOD_REBASING_THRESHOLD</code> value. <p>Default value: <code>FP_CACHE_PERIOD_REBASING_THRESHOLD = 64</code></p> <p>Possible values: 0 (disabled) to 256</p> <p>See "About MSDP storage rebasing" on page 223.</p>

Table 5-27 pd.conf file parameters (continued)

Parameter	Description
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"> ■ The container has been rebased within the last three months. ■ For that backup, the data segments in the container consume less space than the FP_CACHE_REBASING_THRESHOLD value. <p>Default value:</p> <ul style="list-style-type: none"> ■ New installations: FP_CACHE_REBASING_THRESHOLD = 16. ■ System default: If this parameter is not defined in the pd.conf file, NetBackup uses a value of 4. ■ Upgrades: NetBackup uses the system default value of 4. <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 "About MSDP storage rebasing" on page 223.</p>
LOCAL_SETTINGS	<p>Specifies whether to allow the pd.conf settings of the deduplication storage server to override the settings in the local pd.conf file.</p> <p>Set this value to 1 if you use local SEGKSIZE, MINFILE_KSIZE, MATCH_PDRO, and DONT_SEGMENT_TYPES settings.</p> <p>Default value: LOCAL_SETTINGS = 0</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: LOGLEVEL = 0</p> <p>Possible values: An integer, 0 to 10 inclusive</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>

Table 5-27 pd.conf file parameters (continued)

Parameter	Description
MATCH_PDRO	<p>The MATCH_PDRO setting is not included with the default MSDP pd.conf file. This option applies to the NetBackup PureDisk Deduplication Option (PDDO) agent, which is package and installed with PureDisk. The PDDO backs up the data to the standalone PureDisk environment. If you use PDDO in your environment, this option may be in your pd.conf file.</p> <p>Specifies that the PureDisk Deduplication Option should use the same dynamic segmentation deduplication algorithm as the PureDisk Remote Office Agent (PDRO). The following describes the behavior:</p> <ul style="list-style-type: none"> ■ If set to 0 (disabled), the PDDO agent uses a static 128-KB segmentation algorithm for deduplication. ■ If set to 1 (enabled), the PDDO agent uses the same dynamic segmentation deduplication algorithm as the PureDisk Remote Office Agent (PDRO). <p>The data that the PDDO agent backs up is deduplicated globally with the data backed up by the PureDisk environment.</p> <p>Default value: MATCH_PDRO = 1</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>Warning: Symantec recommends that you do not change the MATCH_PDRO value. Changing the default PDDO segmentation algorithm most likely causes the next set of backups to not deduplicate with the existing storage pool data.</p> <p>More information about MATCH_PDRO is available in the following Symantec tech note: http://www.symantec.com/docs/TECH159417</p>
MAX_IMG_MBSIZE	<p>The maximum backup image fragment size in megabytes.</p> <p>Default value: MAX_IMG_MBSIZE = 51200</p> <p>Possible values: 0 to 51,200, in MBs</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>
MAX_LOG_MBSIZE	<p>The maximum size of the log file in megabytes. NetBackup creates a new log file when the log file reaches this limit. NetBackup prepends the date and the ordinal number beginning with 0 to each log file, such as 120131_0_pdplugin.log, 120131_1_pdplugin.log, and so on.</p> <p>Default value: MAX_LOG_MBSIZE = 500</p> <p>Possible values: 0 to 50,000, in MBs</p>

Table 5-27 pd.conf file parameters (continued)

Parameter	Description
META_SEGKSIZE	<p>The segment size for metadata streams</p> <p>Default value: META_SEGKSIZE = 16384</p> <p>Possible values: 32-16384, multiples of 32</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>
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 MaxConcurrentSessions parameter in the mtstrm.conf file controls the number of backups the Multi-Threaded Agent processes concurrently. If you specify more clients than the MaxConcurrentSessions value, some of the clients may use single-threaded processing.</p> <p>See "MSDP mtstrm.conf file parameters" on page 62.</p> <p>The format is a comma-separated list of the clients, case insensitive (for example, MTSTRM_BACKUP_CLIENTS = client1,client2,client3).</p> <p>Default value: MTSTRM_BACKUP_CLIENTS = (no default value)</p> <p>Possible values: comma separated client names</p> <p>See "About the MSDP Deduplication Multi-Threaded Agent" on page 60.</p>

Table 5-27 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"> ■ The hardware concurrency value is one of the following: <ul style="list-style-type: none"> ■ For media servers, half of the host's hardware concurrency is used for the hardware concurrency value in the algorithm. ■ For clients, all of the host's hardware concurrency is used for the hardware concurrency value in the algorithm. ■ The hardware concurrency threshold value to enable multithreading is one of the following: <ul style="list-style-type: none"> ■ Windows and Linux: The threshold value is 2. ■ Solaris: The threshold value is 4. <p>The following examples may be helpful:</p> <ul style="list-style-type: none"> ■ 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>). ■ 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>). <p>See "About the MSDP Deduplication Multi-Threaded Agent" on page 60.</p>

Table 5-27 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 “MSDP mtstrm.conf file parameters” on page 62.</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 “About the MSDP Deduplication Multi-Threaded Agent” on page 60.</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 “About the MSDP Deduplication Multi-Threaded Agent” on page 60.</p>
OPTDUP_BANDWIDTH	<p>Determines the bandwidth that is allowed for each optimized duplication and Auto Image Replication stream on a deduplication server.</p> <p><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 “About configuring MSDP optimized duplication and replication bandwidth” on page 119.</p>

Table 5-27 pd.conf file parameters (continued)

Parameter	Description
OPTDUP_COMPRESSION	<p>Specifies whether to compress optimized duplication data. By default, files are compressed. To disable compression, change the value to 0.</p> <p>Default value: OPTDUP_COMPRESSION = 1</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About MSDP compression” on page 33.</p>
OPTDUP_ENCRYPTION	<p>Specifies whether to encrypt the optimized duplication data. By default, files are not encrypted. If you want encryption, change the value to 1.</p> <p>If you set this parameter to 1 on all hosts, the data is encrypted during transfer and on the storage.</p> <p>Default value: OPTDUP_ENCRYPTION = 0</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About MSDP encryption” on page 34.</p>
OPTDUP_TIMEOUT	<p>Specifies the number of minutes before the optimized duplication times out.</p> <p>Default value: OPTDUP_TIMEOUT = 720</p> <p>Possible values: The value, expressed in minutes</p>
PREFERRED_EXT_SEGKSIZE	<p>Specifies the file extensions and the preferred segment sizes in KB for specific file types. File extensions are case sensitive. The following describe the default values: <code>edb</code> are Microsoft Exchange files; <code>mdf</code> are Microsoft SQL master database files, <code>ndf</code> are Microsoft SQL secondary data files, and <code>segsize64k</code> are Microsoft SQL streams.</p> <p>Default value: PREFERRED_EXT_SEGKSIZE = <code>edb:32,mdf:64,ndf:64,segsize64k:64</code></p> <p>Possible values: <code>file_extension:segment_size_in_KBs</code> pairs, separated by commas.</p> <p>See also SEGKSIZE.</p>
PREFETCH_SIZE	<p>The size in bytes to use for the data buffer for restore operations.</p> <p>Default value: PREFETCH_SIZE = 33554432</p> <p>Possible values: 0 to the computer’s memory limit</p> <p>Note: Change this value only when directed to do so by a Symantec representative.</p>

Table 5-27 pd.conf file parameters (continued)

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

See [“About the MSDP pd.conf configuration file”](#) on page 156.

See [“Editing the MSDP pd.conf file”](#) on page 156.

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 "spallogin" " " 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 172.

See [“Editing an MSDP storage server configuration file”](#) on page 172.

See [“Setting the MSDP storage server configuration”](#) on page 174.

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 171.

See [“Recovering from an MSDP storage server disk failure”](#) on page 231.

See [“Recovering from an MSDP storage server failure”](#) on page 233.

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 5-28](#) shows the configuration lines that are required.

Table 5-28 Required lines for a recovery file

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

See [“About saving the MSDP storage server configuration”](#) on page 171.

See [“Recovering from an MSDP storage server disk failure”](#) on page 231.

See [“Recovering from an MSDP storage server failure”](#) on page 233.

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 172.

- 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 5-28](#)) 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 231.

See [“Recovering from an MSDP storage server failure”](#) on page 233.

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

See [“About saving the MSDP storage server configuration”](#) on page 171.

See [“Saving the MSDP storage server configuration”](#) on page 172.

See [“Editing an MSDP storage server configuration file”](#) on page 172.

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

To set the storage server configuration

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

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

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

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

About the 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.symantecs.org`, the configuration file name is

`DedupeServer.symantecs.org.cfg`.

The following is the location of the file:

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

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

Deleting 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 175.

To delete the host configuration file

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

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

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

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

`DedupeServer.symantecs.org.cfg`

Resetting the 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 198.

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.

Migrating to MSDP storage

This chapter includes the following topics:

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

Replacing the PureDisk Deduplication Option with MSDP on the same host

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

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

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

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

Table 6-1 Replacing a PDDO host with an MSDP host

Step	Task	Procedure
Step 1	Deactivate all backup policies that use the host	Deactivate the policies to ensure that no activity occurs on the host. See the NetBackup Administrator's Guide, Volume I .
Step 2	Upgrade the media server to 7.0 or later	If the media server runs a version of NetBackup earlier than 7.0, upgrade that server to NetBackup 7.0 or later. See the NetBackup Installation Guide for UNIX and Windows .
Step 3	Configure the host	In the Storage Server Configuration Wizard , select PureDisk Deduplication Pool and enter the name of the Storage Pool Authority. See "Configuring a storage server for a Media Server Deduplication Pool" on page 72.
Step 4	Activate your backup policies	See the NetBackup Administrator's Guide, Volume I .

Migrating from PureDisk to the NetBackup MSDP

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

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

Table 6-2 To migrate from PureDisk to NetBackup MSDP

Step	Task	Procedure
Step 1	Install and configure NetBackup	See the NetBackup Installation Guide for UNIX and Windows .
Step 2	Configure NetBackup deduplication	See "Configuring MSDP deduplication" on page 56.
Step 3	Redirect your backup jobs	Redirect your backup jobs to the NetBackup media server deduplication pool. See the NetBackup Administrator's Guide, Volume I .

Table 6-2 To migrate from PureDisk to NetBackup MSDP (*continued*)

Step	Task	Procedure
Step 4	Uninstall PureDisk	Optionally, after the PureDisk backup images expire, uninstall PureDisk. Alternatively, you can continue to use the PureDisk storage as a destination for NetBackup backups. See your NetBackup PureDisk documentation.

See [“Migrating from another storage type to MSDP”](#) on page 179.

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

Table 6-3 Migrating to NetBackup MSDP

Step	Task	Procedure
Step 1	Configure NetBackup deduplication	See “Configuring MSDP deduplication” on page 56.
Step 2	Redirect your backup jobs	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. See the NetBackup Administrator's Guide, Volume I .

Table 6-3 Migrating to NetBackup MSDP (*continued*)

Step	Task	Procedure
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 178.

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.

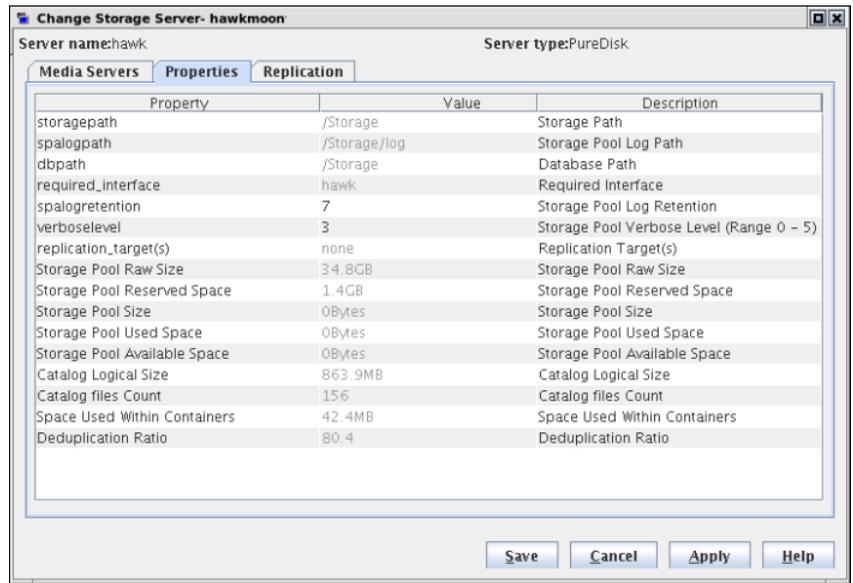
The following methods show the MSDP deduplication rate:

- [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.

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.

Many factors affect deduplication performance.

See [“About MSDP performance”](#) on page 36.

Viewing MSDP job details

Use the NetBackup Activity Monitor 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 183.

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.symantecs.org deduplicated the data (the **dedup** field shows the deduplication rate):

```
03/05/2013 10:37:06 - Info MSDP_Server.symantecs.org (pid=17498)
StorageServer=PureDisk:MSDP_Server.symantecs.org; Report=PDDO Stats
(multi-threaded stream used) for (MSDP_Server.symantecs.org):
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.symantecs.org(pid=2220)
Using OpenStorage client direct to backup from client
Client_B.symantecs.org to MSDP_Server.symantecs.org
1/8/2013 11:58:09 PM - Info MSDP_Server.symantecs.org(pid=2220)
StorageServer=PureDisk:MSDP_Server.symantecs.org; Report=PDDO
```

```

Stats for (MSDP_Server.symantecs.org: 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.symantecs.org(pid=2220)
Using the media server to write NBU data for backup
Client_B_1254987197.symantecs.org to MSDP_Server.symantecs.org
1/8/2013 11:58:19 PM - Info MSDP_Server.symantecs.org(pid=2220)
StorageServer=PureDisk:MSDP_Server.symantecs.org; Report=PDDO
Stats for (MSDP_Server.symantecs.org: scanned: 17161 KB,
CR sent: 17170 KB, dedup: 0.0%, cache hits: 0 (0.0%)

```

Field descriptions

[Table 7-1](#) describes the deduplication activity fields.

Table 7-1 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 cache hits output is not included for the jobs that run on the storage server.</p> <p>See "MSDP pd.conf file parameters" on page 157.</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"> ■ On a NetBackup client that deduplicates its own data (client-side deduplication). ■ On a fingerprinting media server that deduplicates the data. The plug-in on the fingerprinting server sends the data to the storage server, which writes it to a Media Server Deduplication Pool. ■ On a media server that then sends it to a PureDisk environment for storage. (In NetBackup, a PureDisk Storage Pool represents the storage of a PureDisk environment.)
CR sent over FC	<p>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.</p>
dedup	<p>The percentage of data that was stored already. That data is not stored again.</p>
multi-threaded stream used	<p>Indicates that the Deduplication Multi-Threaded Agent processed the backup.</p> <p>See "About the MSDP Deduplication Multi-Threaded Agent" on page 60.</p>

Table 7-1 MSDP activity field descriptions (*continued*)

Field	Description
PDDO stats	Indicates that the job details are for storage on the following destinations: <ul style="list-style-type: none"> ■ Media Server Deduplication Pool ■ PureDisk storage, exposed to NetBackup through a PureDisk Deduplication Pool or the PureDisk Deduplication Option
rebased	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 “About MSDP storage rebasing” on page 223.
scanned	The amount of data that the deduplication plug-in scanned.
Using OpenStorage client direct to restore...	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 7-2](#) describes the options for monitoring capacity and usage.

Table 7-2 Capacity and usage reporting

Option	Description
Change Storage Server dialog box	<p>The Change Storage Server dialog box Properties tab displays storage capacity and usage. It also displays the global deduplication ratio.</p> <p>This dialog box displays the most current capacity usage that is available in the NetBackup Administration Console.</p> <p>You can see an example of the dialog box in a different topic. See “Monitoring the MSDP deduplication rate” on page 181.</p>
Disk Pools window	<p>The Disk Pools 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 Change Storage Server dialog box.</p> <p>To display the window, expand Media and Device Management > Devices > Disk Pools.</p>
View container command	<p>A command that is installed with NetBackup provides a view of storage capacity and usage within the deduplication container files.</p> <p>See “About MSDP container files” on page 187.</p> <p>See “Viewing storage usage within MSDP container files” on page 187.</p>
License Keys dialog box	<p>The summary of active capacity-based license features in the NetBackup License Keys dialog box. The summary displays the storage capacity for which you are licensed and the capacity used. It does not display the amount of physical storage space.</p> <p>On the Help menu in the NetBackup Administration Console, select License Keys.</p>
Disk Pool Status report	<p>The Disk Pool Status report displays the state of the disk pool and usage information.</p> <p>See “Viewing MSDP disk reports” on page 189.</p>

Table 7-2 Capacity and usage reporting (*continued*)

Option	Description
Disk Logs report	<p>The Disk Logs report displays event and message information. A useful event for monitoring capacity is event 1044; the following is the description of the event in the Disk Logs report:</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 “Viewing MSDP disk reports” on page 189.</p> <p>See “MSDP event codes and messages” on page 273.</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 “Determining the MSDP disk volume state” on page 213.</p>
NetBackup OpsCenter	<p>The NetBackup OpsCenter also provides information about storage capacity and usage.</p> <p>See the NetBackup OpsCenter Administrator's Guide.</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

A NetBackup command reports on storage usage within containers.

The following is the pathname of the command:

- UNIX and Linux: `/usr/opensv/pdde/pdcr/bin/crcontrol`.

- Windows: `install_path\Veritas\pdde\Crcontrol.exe`.

To view storage usage within 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 on a Windows deduplication storage server.

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

***** 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 before the space is reserved for metadata.
Size	Raw minus Reserved space.
Used	The file system used space plus the file system reserved space. NetBackup obtains the file system used space from the operating system. 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 reports that space as <code>Used</code> space. Unlike the <code>crcontrol</code> command, some operating system tools report root reserved space as usable space.
Avail	Size minus Used.
Use%	Used divided by 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 7-3](#) describes the disk reports available.

Table 7-3 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 “About MSDP data integrity checking” on page 217.</p> <p>Either <code>PureDisk</code> or <code>Symantec Deduplication Engine</code> in the description identifies a deduplication message. The identifiers are generic because the deduplication engine does not know which application consumes its resources. NetBackup, Symantec Backup Exec, and NetBackup PureDisk are Symantec applications that use deduplication.</p>
Disk Storage Unit 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 Media and Device Management > Devices > Disk Pools.</p> <p>Multiple storage units can point to the same disk pool. When the report query is by storage unit, the report counts the capacity of disk pool storage multiple times.</p>
Disk Pool Status	<p>The Disk Pool Status report displays the state of disk pool and usage information.</p>

To view disk reports

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

About monitoring MSDP processes

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

See [“MSDP server components”](#) on page 237.

Table 7-4 Where to monitor the main MSDP processes

What	Where to monitor it
NetBackup Deduplication Engine	<p>On Windows systems, in the NetBackup Administration Console Activity Monitor Services tab.</p> <p>On UNIX, the NetBackup Deduplication Engine appears as <code>spold</code> in the Administration Console Activity Monitor Daemons tab.</p> <p>The NetBackup <code>bpps</code> command also shows the <code>spold</code> process.</p>
NetBackup Deduplication Manager	<p>On Windows systems, NetBackup Deduplication Manager in the Activity Monitor Services tab.</p> <p>On UNIX, the NetBackup Deduplication Manager appears as <code>spad</code> in the Administration Console Activity Monitor Daemons tab.</p> <p>The NetBackup <code>bpps</code> command also shows the <code>spad</code> process.</p>

Reporting on Auto Image Replication jobs

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

Table 7-5 Auto Image Replication jobs in the Activity Monitor

Job type	Description
Replication	<p>The job that replicates a backup image to a target master displays in the Activity Monitor as a Replication job. The Target Master label displays in the Storage Unit column for this type of job.</p> <p>Similar to other Replication jobs, the job that replicates images to a target master can work on multiple backup images in one instance.</p> <p>The detailed status for this job contains a list of the backup IDs that were replicated.</p>
Import	<p>The job that imports a backup copy into the target master domain displays in the Activity Monitor as an Import job. An Import job can import multiple copies in one instance. The detailed status for an Import job contains a list of processed backup IDs and a list of failed backup IDs.</p> <p>Note: 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 Any 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 Import job fails and NetBackup does not attempt to import the image again.</p> <p>Failed Import jobs fail with a status 191 and appear in the Problems report when run on the target master server.</p> <p>The image is expired and deleted during an Image Cleanup job. Note that the originating domain (Domain 1) does not track failed imports.</p>

Managing deduplication

This chapter includes the following topics:

- [Managing 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](#)
- [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 193.
- See [“Determining the MSDP storage server state”](#) on page 193.
- See [“Viewing MSDP storage server attributes”](#) on page 194.
- See [“Setting MSDP storage server attributes”](#) on page 195.
- See [“Changing MSDP storage server properties”](#) on page 196.
- See [“Clearing MSDP storage server attributes”](#) on page 197.
- See [“About changing the MSDP storage server name or storage path”](#) on page 198.
- See [“Changing the MSDP storage server name or storage path”](#) on page 198.
- See [“Removing an MSDP load balancing server”](#) on page 199.
- See [“Deleting an MSDP storage server”](#) on page 201.
- See [“Deleting the MSDP storage server configuration”](#) on page 201.

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 206.

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:

`-storage_server` The name of the storage server.
`storage_server`

`-setattribute` The *attribute* is the name of the argument that represents the
`attribute` new functionality.

For example, **OptimizedImage** specifies that the environment supports the optimized synthetic backup method.

The following is the path to the `nbdevconfig` command:

- UNIX: `/usr/opensv/netbackup/bin/admincmd`
 - Windows: `install_path\NetBackup\bin\admincmd`
- 2 To verify, view the storage server attributes.
- See [“Viewing MSDP storage server attributes”](#) on page 194.
- See [“About MSDP optimized synthetic backups”](#) on page 35.

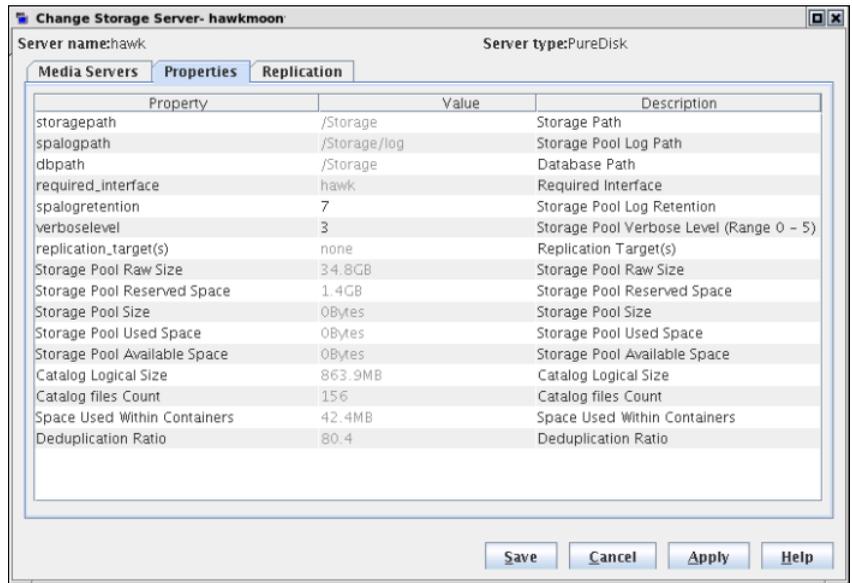
Changing MSDP storage server properties

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**.

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



- For the property to change, select the value in the **Value** column.
- Change the value.
- 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 198.
- 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 198.
- 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 233.

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 198.

Table 8-1 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 NetBackup Administrator's Guide, Volume I .

Table 8-1 Changing the storage server name or storage path (*continued*)

Step	Task	Procedure
Step 2	Expire the backup images	<p>Expire all backup images that reside on the deduplication disk storage.</p> <p>Warning: Do not delete the images. They are imported back into NetBackup later in this process.</p> <p>If you use the <code>bpxpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter.</p> <p>See the NetBackup Administrator's Guide, Volume I.</p>
Step 3	Delete the storage units that use the disk pool	See the NetBackup Administrator's Guide, Volume I .
Step 4	Delete the disk pool	See "Deleting a Media Server Deduplication Pool" on page 215.
Step 5	Delete the deduplication storage server	See "Deleting an MSDP storage server" on page 201.
Step 6	Delete the configuration	<p>Delete the deduplication configuration.</p> <p>See "Deleting the MSDP storage server configuration" on page 201.</p>
Step 7	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 175.</p>
Step 8	Change the storage server name or the storage location	<p>See the computer or the storage vendor's documentation.</p> <p>See "Use fully qualified domain names" on page 38.</p> <p>See "About the MSDP storage paths" on page 50.</p>
Step 9	Reconfigure the storage server	<p>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.</p> <p>See "Configuring MSDP deduplication" on page 56.</p>
Step 10	Import the backup images	See the NetBackup Administrator's Guide, Volume I .

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 25.

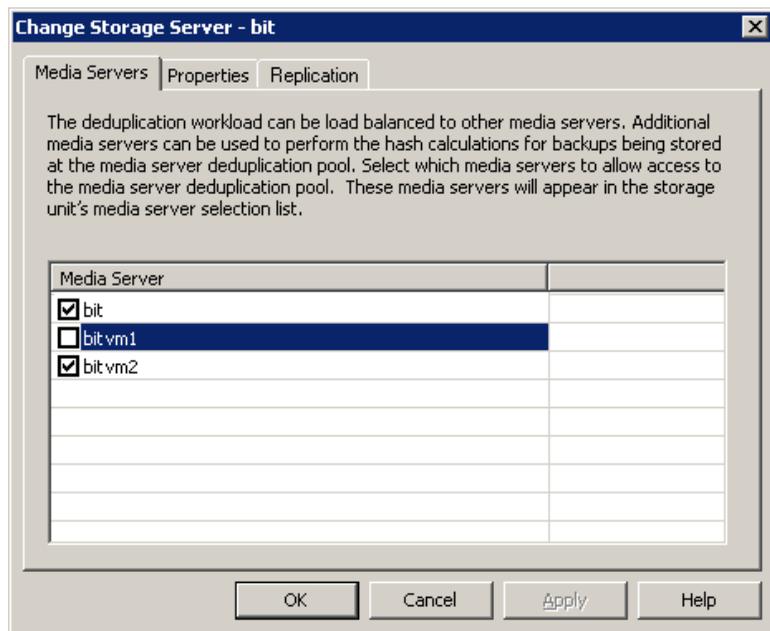
After you remove the load balancing server, restart the NetBackup Enterprise Media Manager service. The NetBackup disk polling service may try to use the removed server to query for disk status. Because the server is no longer a load balancing server, it cannot query the disk storage. Consequently, NetBackup may mark the disk volume as DOWN. When the EMM service restarts, it chooses a different deduplication server to monitor the disk storage.

If the host failed and is unavailable, you can use the `tpconfig` device configuration utility in menu mode to delete the server. However, you must run the `tpconfig` utility on a UNIX or Linux NetBackup server.

For procedures, see the [NetBackup Administrator's Guide, Volume II](#).

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 198.

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 198.

See [“Removing MSDP”](#) on page 235.

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 Run one of the following scripts, depending on your operating system:

UNIX:

```
/usr/opensv/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 203.

See [“Adding NetBackup Deduplication Engine credentials”](#) on page 203.

See [“Changing NetBackup Deduplication Engine credentials”](#) on page 204.

See [“Deleting credentials from a load balancing server”](#) on page 204.

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 155.

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/openv/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 NetBackup Deduplication Engine credentials”](#) on page 31.

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 199.

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 205.

See [“Determining the Media Server Deduplication Pool state”](#) on page 205.

See [“Changing a Media Server Deduplication Pool state”](#) on page 205.

See [“Viewing Media Server Deduplication Pool attributes”](#) on page 206.

See [“Setting a Media Server Deduplication Pool attribute”](#) on page 206.

See [“Changing a Media Server Deduplication Pool properties”](#) on page 207.

See [“Clearing a Media Server Deduplication Pool attribute”](#) on page 212.

See [“Determining the MSDP disk volume state”](#) on page 213.

See [“Changing the MSDP disk volume state”](#) on page 213.

See [“Inventorying a NetBackup disk pool”](#) on page 214.

See [“Deleting a Media Server Deduplication Pool”](#) on page 215.

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 NetBackup Media Server Deduplication Pools”](#) on page 83.

See [“Determining the Media Server Deduplication Pool state”](#) on page 205.

See [“Media Server Deduplication Pool properties”](#) on page 91.

See [“Configuring a Media Server Deduplication Pool”](#) on page 84.

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.symantecs.org (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 195.

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 206.

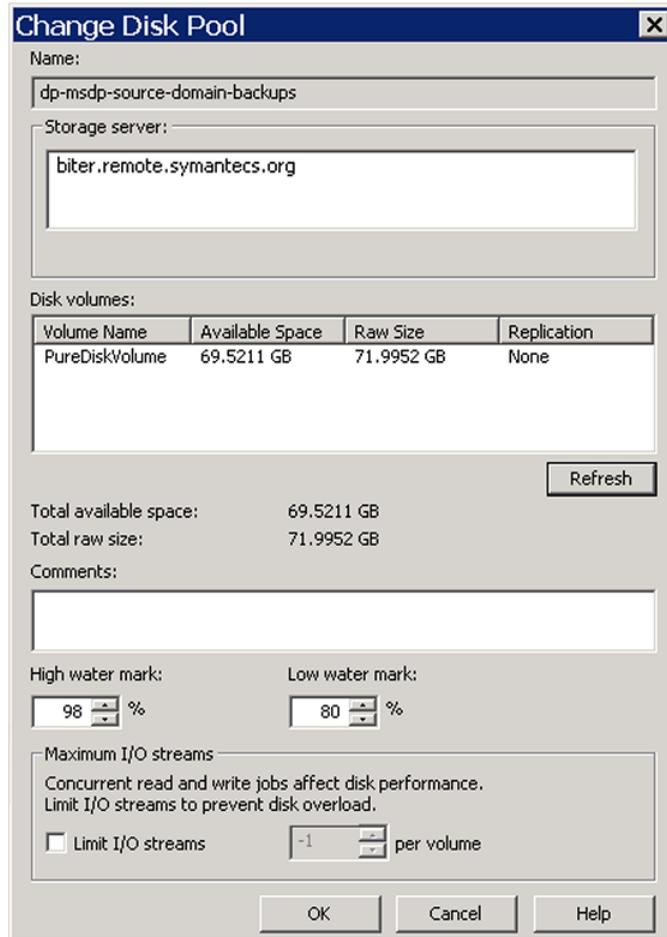
Changing a Media Server Deduplication Pool properties

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**.



- 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 209.

- 5 Change the other properties as necessary.

See [“Media Server Deduplication Pool properties”](#) on page 91.

- 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 8-2 Refresh outcomes

Outcome	Description
No changes are discovered.	No changes are required.
NetBackup discovers the new volumes that you can add to the disk pool.	The new volumes appear in the Change Disk Pool dialog box. Text in the dialog box changes to indicate that you can add the new volumes to the disk pool.

Table 8-2 Refresh outcomes (*continued*)

Outcome	Description
<p>The replication properties of all of the volumes changed, but they are still consistent.</p>	<p>A Disk Pool Configuration Alert pop-up box notifies you that the properties of all of the volumes in the disk pool changed, but they are all the same (homogeneous).</p>  <p>You must click OK in the alert box, after which the disk pool properties in the Change Disk Pool dialog box are updated to match the new volume properties.</p> <p>If new volumes are available that match the new properties, NetBackup displays those volumes in the Change Disk Pool dialog box. You can add those new volumes to the disk pool.</p> <p>In the Change Disk Pool dialog box, select one of the following two choices:</p> <ul style="list-style-type: none"> ■ OK. To accept the disk pool changes, click OK in the Change Disk Pool dialog box. NetBackup saves the new properties of the disk pool. NetBackup can use the disk pool, but it may no longer match the intended purpose of the storage unit or storage lifecycle policy. Change the storage lifecycle policy definitions to ensure that the replication operations use the correct source and target disk pools, storage units, and storage unit groups. Alternatively, work with your storage administrator to change the volume properties back to their original values. ■ Cancel. To discard the changes, click Cancel in the Change Disk Pool dialog box. NetBackup does not save the new disk pool properties. NetBackup can use the disk pool, but it may no longer match the intended use of the storage unit or storage lifecycle policy.

Table 8-2 Refresh outcomes (*continued*)

Outcome	Description
<p>The replication properties of the volumes changed, and they are now inconsistent.</p>	<p>A Disk Pool Configuration Error pop-up box notifies you that the replication properties of some of the volumes in the disk pool changed. The properties of the volumes in the disk pool are not homogeneous.</p>  <p>You must click OK in the alert box.</p> <p>In the Change Disk Pool dialog box, the properties of the disk pool are unchanged, and you cannot select them (that is, they are dimmed). However, the properties of the individual volumes are updated.</p> <p>Because the volume properties are not homogeneous, NetBackup cannot use the disk pool until the storage configuration is fixed.</p> <p>NetBackup does not display new volumes (if available) because the volumes already in the disk pool are not homogeneous.</p> <p>To determine what has changed, compare the disk pool properties to the volume properties.</p> <p>See “Viewing the replication topology for Auto Image Replication” on page 128.</p> <p>Work with your storage administrator to change the volume properties back to their original values.</p> <p>The disk pool remains unusable until the properties of the volumes in the disk pool are homogenous.</p> <p>In the Change Disk Pool dialog box, click OK or Cancel to exit the Change Disk Pool dialog box.</p>

Table 8-2 Refresh outcomes (*continued*)

Outcome	Description
<p>NetBackup cannot find a volume or volumes that were in the disk pool.</p>	<p>A Disk Pool Configuration Alert pop-up box notifies you that an existing volume or volumes was deleted from the storage device:</p>  <p>NetBackup can use the disk pool, but data may be lost.</p> <p>To protect against accidental data loss, NetBackup does not allow volumes to be deleted from a disk pool.</p> <p>To continue to use the disk pool, do the following:</p> <ul style="list-style-type: none"> ■ Use the <code>bpimmedia</code> command or the Images on Disk report to display the images on the specific volume. ■ Expire the images on the volume. ■ Use the <code>nbdevconfig</code> command to set the volume state to DOWN so NetBackup does not try to use it.

Clearing 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/opensv/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/opensv/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 266.

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 225.

To delete backup images manually

- 1 Expire all of the images by using the `bpeupdate` command and the `-notimmediate` option. The `-notimmediate` option prevents `bpeupdate` from calling the `nbdelete` command, which deletes the image.

Without this option, `bpeupdate` 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 `storaged.log` file.

See [“About MSDP log files”](#) on page 253.

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 216.

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

Processing the MSDP transaction queue manually

Usually, you should not need to run the deduplication database transaction queue processes manually. However, you can do so.

See [“About MSDP queue processing”](#) on page 216.

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 189.

The data integrity message code is 1057.

See [“MSDP event codes and messages”](#) on page 273.

NetBackup writes the integrity checking activity messages to the NetBackup Deduplication Engine `storaged.log` file.

See [“About MSDP log files”](#) on page 253.

You can configure some of the data integrity checking behaviors.

See [“Configuring MSDP data integrity checking behavior”](#) on page 218.

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 218.
- Edit configuration file parameters.
See [“To configure data integrity checking behavior by editing the configuration files”](#) on page 219.

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 217.

See [“MSDP data integrity checking configuration parameters”](#) on page 220.

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	UNIX: <code>/usr/openv/pdde/pdcr/bin/pddecfg -a enabledataintegritycheck</code>
--------	---

Disable UNIX: `/usr/opensv/pdde/pdcr/bin/pddecfg -a disabledataintegritycheck`

Get the status UNIX: `/usr/opensv/pdde/pdcr/bin/pddecfg -a getdataintegritycheck`

- **Cyclic redundancy checking. Use the following commands to configure behavior:**

Enable CRC does not run if queue processing is active or during disk read or write operations.

UNIX: `/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckon`

Windows: `install_path\Veritas\pdde\Crcontrol.exe --crccheckon`

Disable UNIX: `/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckoff`

Windows: `install_path\Veritas\pdde\Crcontrol.exe --crccheckoff`

Enable fast checking Fast check CRC mode begins the check from container 64 and does not sleep between checking containers.

When the fast CRC ends, CRC behavior reverts to the behavior before fast checking was invoked.

UNIX: `/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckrestart`

Windows: `install_path\Veritas\pdde\Crcontrol.exe --crccheckrestart`

Get the status UNIX: `/usr/opensv/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 220.
 - 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 8-3](#).
 See “[About the MSDP contentrouter.cfg file](#)” on page 170.
- The `spa.cfg` file.
 The parameters are described in [Table 8-3](#).

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 217.

Table 8-3 The `contentrouter.cfg` file parameters for data integrity checking

Setting	Default	Description
EnableCRCCheck	true	Enable or disable cyclic redundancy checking (CRC) of the data container files. The possible values are <code>true</code> or <code>false</code> . CRC occurs only when no backup, restore, or queue processing jobs are running.

Table 8-3 The contentrouter.cfg file parameters for data integrity checking (continued)

Setting	Default	Description
CRCCheckSleepSeconds	5	The time in seconds to sleep between checking containers. The longer the sleep interval, the more time it takes to check containers.
CRCCheckBatchNum	40	The number of containers to check each time. The greater the number of containers, the less time it takes to check all containers, but the more system resources it takes.
ShutdownCRWhenError	false	Stops the NetBackup Deduplication Manager when a data loss is discovered. This parameter is reserved for debugging purposes by Symantec Support Representatives. The possible values are <code>true</code> or <code>false</code> .
GarbageCheckRemainDCCount	100	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.

Table 8-4 spa.cfg file parameters for data integrity checking

Setting	Default	Description
EnableDataCheck	true	Enable or disable data consistency checking. The possible values are <code>True</code> or <code>False</code> .
DataCheckDays	14	The number of days to check the data for consistency. The greater the number of days, the fewer the objects that are checked each day. The greater the number of days equals fewer storage server resources consumed each day.
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 “About MSDP log files” on page 253.

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 and duplication jobs to tape. Duplication jobs to tape rehydrate the data into the original NetBackup backup images and write the images to tape for long-term archiving.

In most cases, you should change configuration file options only when directed to do so by Symantec Technical Support.

Table 8-5 Some settings that affect restore operations

Process	Description
Defragmenting the storage	<p>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.</p> <p>If you upgrade from a NetBackup release earlier than 7.5, rebasing may affect your deduplication performance temporarily.</p> <p>See “About MSDP storage rebasing” on page 223.</p>
Prefetching data	<p>You can control the prefetching behavior as follows:</p> <ul style="list-style-type: none"> ■ The <code>PrefetchThreadNum</code> parameter in the <code>contentrouter.cfg</code> file specifies the number of threads to use to preload segments during storage read operations. The default value is 1. Depending on your disks, a value as high as 4 may improve performance. However, Symantec recommends that you test higher values thoroughly to ensure that a value greater than 1 yields better performance. Higher values may actually decrease performance. See “About the MSDP contentrouter.cfg file” on page 170. ■ The <code>PREFETCH_SIZE</code> parameter in the <code>pd.conf</code> file specifies the size in bytes to use for the data buffer for read operations. See “About the MSDP pd.conf configuration file” on page 156. See “MSDP pd.conf file parameters” on page 157.
Decrypting the data on the client rather than the server	<p>The <code>RESTORE_DECRYPT_LOCAL</code> parameter in the <code>pd.conf</code> file specifies on which host to decrypt and decompress the data during restore operations.</p> <p>See “About the MSDP pd.conf configuration file” on page 156.</p> <p>See “MSDP pd.conf file parameters” on page 157.</p>

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 8-6](#) describes the rebasing operations.

Table 8-6 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"> ■ The container has been rebased within the last three months. ■ 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. See “MSDP <code>pd.conf</code> file parameters” on page 157. <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 8-6 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"> ■ The container has not been rebased within the last three months. ■ 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. <p>See “MSDP <code>pd.conf</code> file parameters” on page 157.</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"> ■ <code>RebaseScatterThreshold</code> ■ <code>RebaseMinContainers</code> ■ <code>RebaseMaxPercentage</code> ■ <code>RebaseMaxTime</code> <p>See “MSDP server-side rebasing parameters” on page 224.</p>

MSDP server-side rebasing parameters

[Table 8-7](#) describes the parameters that control server-side rebasing.

See “[About MSDP storage rebasing](#)” on page 223.

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 170.

Table 8-7 The server-side rebasing parameters

Parameter	Description
RebaseMaxPercentage	<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>
RebaseMaxTime	<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>
RebaseMinContainers	<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>
RebaseScatterThreshold	<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 187.

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 217.

See [“Deleting backup images”](#) on page 215.

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](#).
- 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](#).
- 7 Restart the deduplication jobs.

About restoring files at a remote site

If you use optimized duplication to copy images from a local site to a remote site, you can restore from the copies at the remote site to clients at the remote site. To do so, use a server-directed restore or a client-redirected restore, which restores files to a client other than the original client.

Information about how to redirect restores is in a different guide.

See “Managing client restores” in the [NetBackup Administrator's Guide, Volume I](#).

You may have to configure which media server performs the restore. In optimized duplication, the media server that initiates the duplication operation becomes the write host for the new image copies. The write host restores from those image copies. If the write host is at the local site, it restores from those images at the remote site to the alternate client at the remote site. That host reads the image across the WAN and then writes the image back across the WAN to the alternate client. In this case, you can specify that the media server at the remote site as the restore server.

About restoring from a backup at a target master domain

While it is possible to restore a client directly by using the images in the target master domain, do so only in a disaster recovery situation. In this discussion, a disaster recovery situation is one in which the originating domain no longer exists and clients must be recovered from the target domain.

Table 8-8 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 43.

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](#).
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/openv/netbackup/db/config`
Windows: `install_path\veritas\netbackup\db\config`
This global setting always forces restores to the server that did the backup. It applies to all NetBackup restore jobs, not just deduplication restore jobs. If this touch file exists, NetBackup ignores the `FORCE_RESTORE_MEDIA_SERVER` and `FAILOVER_RESTORE_MEDIA_SERVER` settings.
- Always use a different server.
Use NetBackup **Host Properties** to specify a **Media host override** server. See the previous explanation about **Media host override**, except: Specify the different server for the **Restore server**.
- A single restore instance. Use the `bprestore` command with the `-disk_media_server` option.
Restore jobs for each instance of the command use the media server you specify. See [NetBackup Commands Reference Guide](#).

Host replacement, recovery, and uninstallation

This chapter includes the following topics:

- [Replacing the MSDP storage server host computer](#)
- [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 uninstalling MSDP](#)
- [Removing MSDP](#)

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 host include a lease swap or perhaps the current deduplication storage server host does not meet your performance requirements.

Warning: The new host should use the same version of NetBackup and the same operating system and version. If you use a more recent version of NetBackup on the replacement host, ensure that you perform any data conversions that may be required for the newer release.

The new host also must use the same byte order as the old host. If it does not, you cannot access the deduplicated data. In computing, endianness describes the byte order that represents data: big endian and little endian. For example, SPARC processors and Intel processors use different byte orders. Therefore, you cannot replace an Oracle Solaris SPARC host with an Oracle Solaris host that has an Intel processor.

Table 9-1 Replacing an MSDP storage server host computer

Step	Task	Procedure
Step 1	Expire the backup images	Expire all backup images that reside on the deduplication disk storage. Warning: Do not delete the images. They are imported back into NetBackup later in this process. If you use the <code>bpxpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter. See the NetBackup Administrator's Guide, Volume I .
Step 2	Delete the storage units that use the disk pool	See the NetBackup Administrator's Guide, Volume I .
Step 3	Delete the disk pool	See "Deleting a Media Server Deduplication Pool" on page 215.
Step 4	Delete the deduplication storage server	See "Deleting an MSDP storage server" on page 201.
Step 5	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See "Deleting an MSDP host configuration file" on page 175.
Step 6	Delete the credentials on deduplication servers	If you have load balancing servers, delete the NetBackup Deduplication Engine credentials on those media servers. See "Deleting credentials from a load balancing server" on page 204.

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"> ■ You can use the same host name or a different name. ■ You can use the same Storage Path or a different Storage Path. If you use a different Storage Path, you must move the deduplication storage to that new location. ■ If the Database Path on the original host is different that the Storage Path, you can do one of the following: <ul style="list-style-type: none"> ■ You can use the same Database Path. ■ You can use a different Database Path. If you do, you must move the deduplication database to the new location. ■ You do not have to continue to use a different Database Path. You can move the <code>databases</code> directory into the Storage Path and then specify only the Storage Path when you configure the storage server. ■ 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. <p>See "About MSDP storage servers" on page 25.</p> <p>See "About MSDP server requirements" on page 26.</p>
Step 8	Connect the storage to the host	<p>Use the storage path that you configured for this replacement host.</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 NetBackup Installation Guide for UNIX and Windows.</p>
Step 10	Reconfigure deduplication	<p>See "Configuring MSDP deduplication" on page 56.</p>
Step 11	Import the backup images	<p>See the NetBackup Administrator's Guide, Volume I.</p>

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 9-2 Process to recover from media server disk failure

Step	Task	Procedure
Step 1	Replace the disk.	If the disk is a system boot disk, also install the operating system. See the hardware vendor and operating system documentation.
Step 2	Mount the storage.	Ensure that the storage and database are mounted at the same locations. See the storage vendor's documentation.
Step 3	Install and license the NetBackup media server software.	See the NetBackup Installation Guide for UNIX and Windows . See "About the MSDP license key" on page 54.
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 175.
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 204.
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 203.
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 172.
Step 8	Edit the configuration file	See "Editing an MSDP storage server configuration file" on page 172.

Table 9-2 Process to recover from media server disk failure (*continued*)

Step	Task	Procedure
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 174.
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 155.

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.

Warning: The new host should use the same version of NetBackup and the same operating system and version. If you use a more recent version of NetBackup on the replacement host, ensure that you perform any data conversions that may be required for the newer release.

The new host also must use the same byte order as the old host. If it does not, you cannot access the deduplicated data. In computing, endianness describes the byte order that represents data: big endian and little endian. For example, SPARC processors and Intel processors use different byte orders. Therefore, you cannot replace an Oracle Solaris SPARC host with an Oracle Solaris host that has an Intel processor.

Table 9-3 Recover from an MSDP storage server failure

Step	Task	Procedure
Step 1	Expire the backup images	Expire all backup images that reside on the deduplication disk storage. Warning: Do not delete the images. They are imported back into NetBackup later in this process. If you use the <code>bpxpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter. See the NetBackup Administrator's Guide, Volume I .
Step 2	Delete the storage units that use the disk pool	See the NetBackup Administrator's Guide, Volume I .

Table 9-3 Recover from an MSDP storage server failure (*continued*)

Step	Task	Procedure
Step 3	Delete the disk pool	See “Deleting a Media Server Deduplication Pool” on page 215.
Step 4	Delete the deduplication storage server	See “Deleting an MSDP storage server” on page 201.
Step 5	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See “Deleting an MSDP host configuration file” on page 175.
Step 6	Delete the credentials on deduplication servers	If you have load balancing servers, delete the NetBackup Deduplication Engine credentials on those media servers. See “Deleting credentials from a load balancing server” on page 204.
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"> ■ You can use the same host name or a different name. ■ You can use the same Storage Path or a different Storage Path. If you use a different Storage Path, you must move the deduplication storage to that new location. ■ If the Database Path on the original host is different than the Storage Path, you can do one of the following: <ul style="list-style-type: none"> ■ You can use the same Database Path. ■ You can use a different Database Path. If you do, you must move the deduplication database to the new location. ■ You do not have to continue to use a different Database Path. You can move the <code>databases</code> directory into the Storage Path and then specify only the Storage Path when you configure the storage server. ■ 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. See “About MSDP storage servers” on page 25. See “About MSDP server requirements” on page 26.
Step 8	Connect the storage to the host	Use the storage path that you configured for this replacement host. See the computer or the storage vendor’s documentation.
Step 9	Install the NetBackup media server software on the new host	See the NetBackup Installation Guide for UNIX and Windows .

Table 9-3 Recover from an MSDP storage server failure (*continued*)

Step	Task	Procedure
Step 10	Reconfigure deduplication	You must use the same credentials for the NetBackup Deduplication Engine. See “Configuring MSDP deduplication” on page 56.
Step 11	Import the backup images	See the NetBackup Administrator's Guide, Volume I .

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 174.

Symantec recommends that you save your storage server configuration.

See [“Save the MSDP storage server configuration”](#) on page 42.

Information about recovering the master server is available.

See the [NetBackup Troubleshooting Guide](#).

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 198.
- Deactivate deduplication and remove the configuration files and the storage files.
See [“Removing MSDP”](#) on page 235.

Removing 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 9-4 Remove MSDP

Step	Task	Procedure
Step 1	Remove client deduplication	Remove the clients that deduplicate their own data from the client deduplication list. See “Disabling MSDP client-side deduplication for a client” on page 149.
Step 2	Delete the storage units that use the disk pool	See the NetBackup Administrator's Guide, Volume I .
Step 3	Delete the disk pool	See “Deleting a Media Server Deduplication Pool” on page 215.
Step 4	Delete the deduplication storage server	See “Deleting an MSDP storage server” on page 201. Deleting the deduplication storage server does not alter the contents of the storage on physical disk. To protect against inadvertent data loss, NetBackup does not automatically delete the storage when you delete the storage server.
Step 5	Delete the configuration	Delete the deduplication configuration. See “Deleting the MSDP storage server configuration” on page 201.
Step 6	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See “Deleting an MSDP host configuration file” on page 175.
Step 7	Delete the storage directory and the database directory	Delete the storage directory and database directory. (Using a separate database directory was an option when you configured deduplication.) Warning: If you delete the storage directory and valid NetBackup images reside on the deduplication storage, data loss may occur. See the operating system documentation.

Deduplication architecture

This chapter includes the following topics:

- MSDP server components
- Media server deduplication backup process
- MSDP client components
- MSDP client-side deduplication backup process

MSDP server components

Figure 10-1 is a diagram of the storage server components.

Figure 10-1 MSDP server components

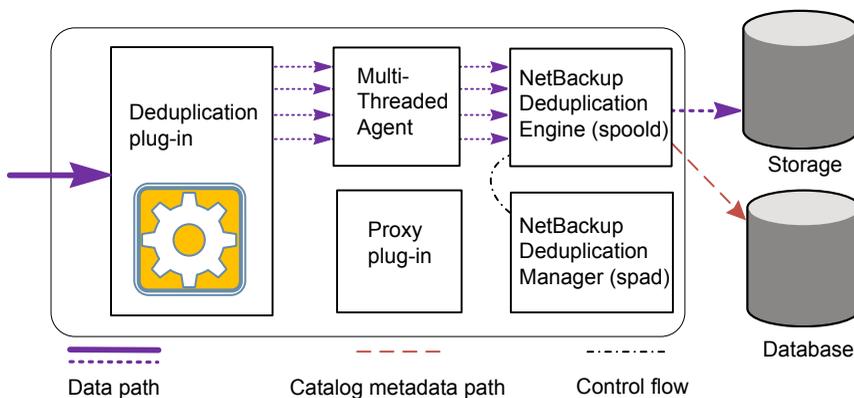


Table 10-1 describes the MSDP storage server components.

Table 10-1 NetBackup MSDP server components

Component	Description
Deduplication plug-in	<p>The deduplication plug-in does the following:</p> <ul style="list-style-type: none"> ■ Separates the file's metadata from the file's content. ■ Deduplicates the content (separates files into segments). <p>The plug-in runs on the deduplication storage server, on load balancing servers, and on the clients that deduplicate their own data.</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 “About the MSDP Deduplication Multi-Threaded Agent” on page 60.</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 Table 10-2.</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>nboostpxy</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 10-2](#) describes the components and functions within the NetBackup Deduplication Engine.

Table 10-2 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"> ■ Provides a thread pool to serve all clients. ■ Maintains a task for each client connection. ■ Manages the mode of the Deduplication Engine based on the operation. Operations are backups, restores, queue processing, and so on.
Data integrity checking	<p>The NetBackup Deduplication Engine checks the integrity of the data and resolves integrity problems.</p> <p>See “About MSDP data integrity checking” on page 217.</p>
Data Store Manager	<p>The Data Store Manager manages all of the data container files. The Data Store Manager is a set of functions and threads that provides the following:</p> <ul style="list-style-type: none"> ■ A transaction mechanism to back up data into the data store. ■ A mechanism to read data from the data store. ■ A transaction mechanism to reclaim space in the data store (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. <p>See “About MSDP container files” on page 187.</p>
Index Cache Manager	<p>The Index Cache Manager manages the fingerprint cache. The cache improves fingerprint lookup speed.</p> <p>See “About the MSDP fingerprint cache” on page 67.</p>
Queue processing	<p>The NetBackup Deduplication Engine processes the transaction queue.</p> <p>See “About MSDP queue processing” on page 216.</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>

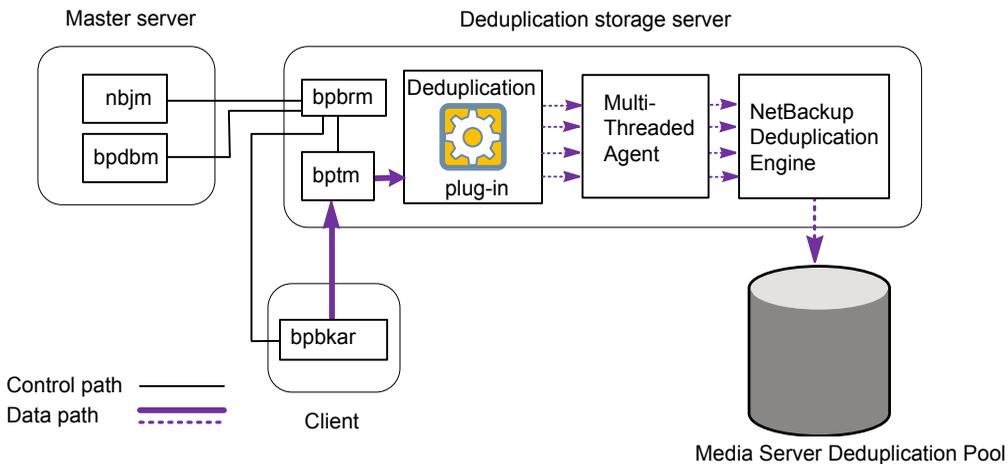
Table 10-2 NetBackup Deduplication Engine components and functions
(continued)

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 10-2](#) diagram shows the backup process when a media server deduplicates the backups. The destination is a **Media Server Deduplication Pool**. A description follows.

Figure 10-2 Media server deduplication process



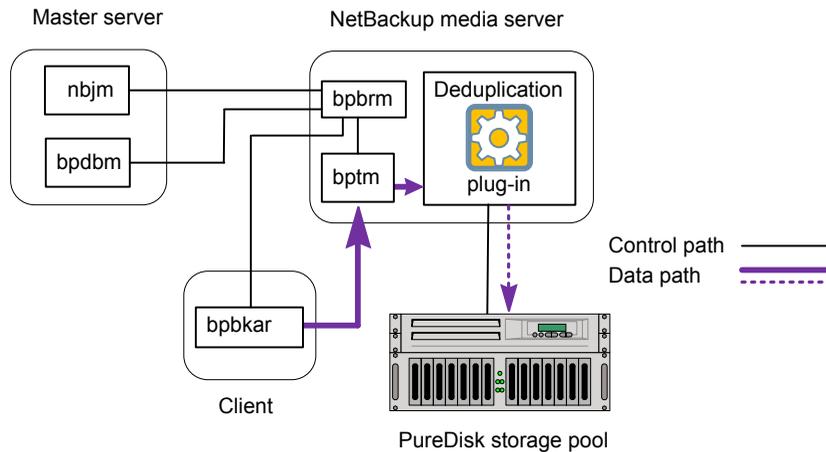
The following list describes the backup process when a media server deduplicates the backups and the destination is a **Media Server Deduplication Pool**:

- The NetBackup Job Manager (`nbjm`) starts the Backup/Restore Manager (`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 `bpdm` 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 10-3](#) diagram shows the backup process when a media server deduplicates the backups. The destination is a **PureDisk Deduplication Pool**. A description follows.

Figure 10-3 Media server deduplication process to a PureDisk storage pool

The following list describes the backup process when a media server deduplicates the backups and the destination is a **PureDisk Deduplication Pool**:

- The NetBackup Job Manager (`nbjm`) starts the Backup/Restore Manager (`bpbrm`) on a media server.
- The Backup/Restore Manager starts the `bptm` process on the media server and the `bpbkar` process on the client).
- The Backup/Archive Manager (`bpbkar`) generates the backup images and moves them to the media server `bptm` process.
The Backup/Archive Manager also sends the information about files within the image to the Backup/Restore Manager (`bpbrm`). The Backup/Restore Manager sends the file information to the `bpdbm` process on the master server for the NetBackup database.
- The `bptm` process moves the data to the deduplication plug-in.
- The deduplication plug-in retrieves a list of 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 10-3](#) describes the client deduplication components.

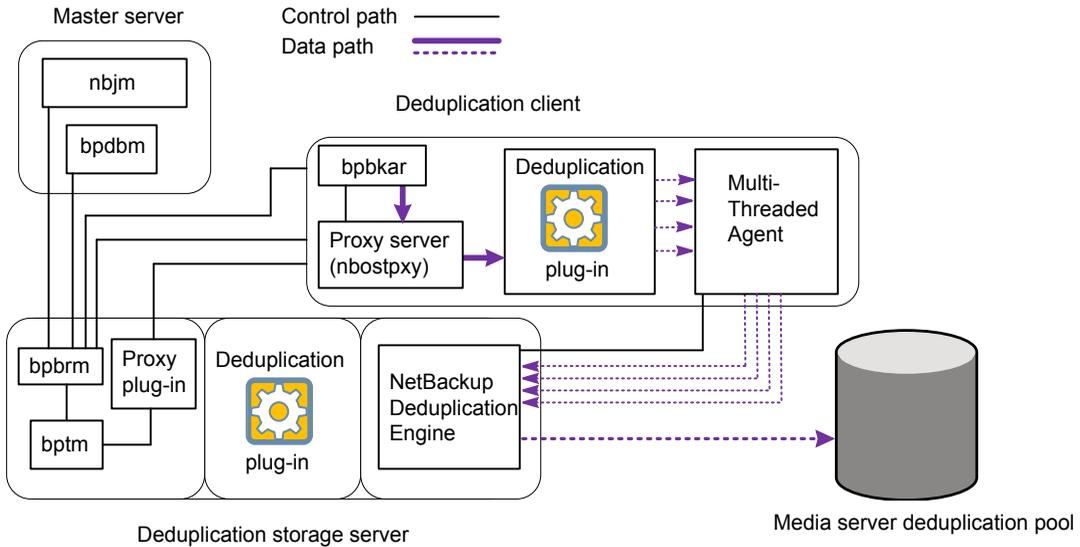
Table 10-3 Client MSDP components

Component	Description
Deduplication plug-in	The deduplication plug-in does the following: <ul style="list-style-type: none">■ Separates the file's metadata from the file's content.■ Deduplicates the content (separates files into segments).
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 deduplicate their own data. See “About the MSDP Deduplication Multi-Threaded Agent” on page 60.
Proxy server	The OpenStorage proxy server (<code>nbostproxy</code>) manages control communication with the proxy plug-in on the storage server.

MSDP client–side deduplication backup process

The [Figure 10-4](#) 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 10-4 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 69.
 - 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 69.

The list of fingerprints is used as a cache so the plug-in does not have to request each fingerprint from the engine.

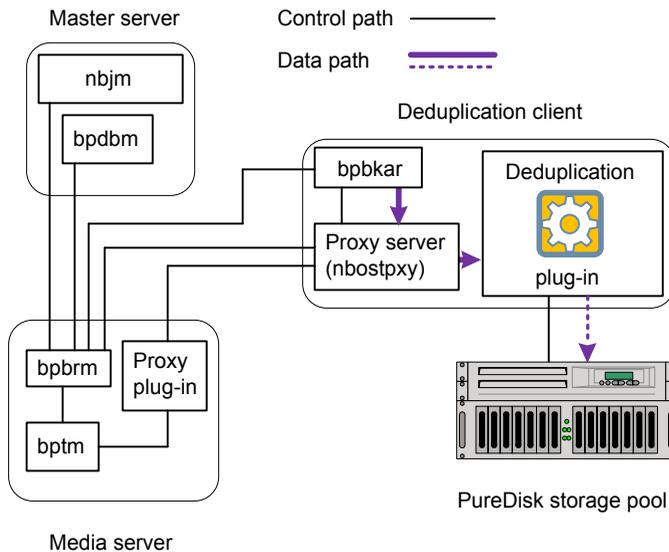
If no fingerprints are loaded into the cache, the deduplication rate may be very low for the backup.

- 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 10-5](#) diagram shows the backup process of a client that deduplicates its own data. The destination is a PureDisk storage pool. A description follows.

Figure 10-5 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 (*bpbrm*) on a media server.
- The Backup / Restore Manager probes the client to determine if it is configured and ready for deduplication.
- If the client is ready, the Backup/Restore Manager starts the following processes: The OpenStorage proxy server (*nbostpxy*) on the client and the data moving processes (*bpbkar* on the client and *bptm* on the media server). NetBackup uses the proxy plug-in on the media server to route control information from *bptm* to *nbostpxy*.
- The Backup/Archive Manager (*bpbkar*) generates the backup images and moves them to the client *nbostpxy* process by shared memory. The Backup/Archive Manager also sends the information about files within the image to the Backup/Restore Manager (*bpbrm*). The Backup/Restore Manager sends the file information to the *bpdbm* process on the master server for the NetBackup database.
- The client *nbostpxy* process moves the data to the deduplication plug-in.

- The deduplication plug-in retrieves a list of fingerprints from the last full backup 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 251.

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:

```
UNIX      /usr/opensv/logs
Windows  install_path\NetBackup\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:

```
vxlogcfg      Modifies the unified logging configuration settings.
               for more information about the vxlogcfg command.

vxlogmgr      Manages the log files that the products that support unified logging
               generate.
               for more information about the vxlogmgr command.

vxlogview     Displays the logs that unified logging generates.
               See "Examples of using vxlogview to view unified logs" on page 250.
               for more information about the vxlogview command.
```

See the [NetBackup Commands Reference Guide](#) for a complete description about these commands.

These commands are located in the following directory:

```
UNIX      /usr/opensv/netbackup/bin
Windows  install_path\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\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 11-1 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 11-1 Example uses of the vxlogview command (*continued*)

Item	Example
<p>Display results faster</p>	<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>Note: 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>
<p>Search for a job ID</p>	<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>

A complete description of `vxlogview` is in the [NetBackup Commands Reference Guide](#).

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 248.

To enable legacy debug logging on NetBackup servers, you must first create the appropriate directories for each process.

UNIX /usr/opensv/netbackup/logs
 /usr/opensv/volmgr/debug

Windows install_path\NetBackup\logs
 install_path\Volmgr\debug

After the directories are created, NetBackup creates log files in the directory that is 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`.

As an alternative, you can stop and restart `nbemm` after creating the following directory:

UNIX /usr/opensv/volmgr/debug/reqlib

Windows install_path\Volmgr\debug\reqlib\

Tables are available that list the log directories that you must create.

Note: On a Windows server, you can create the debug log directories at once, under `install_path\NetBackup\Logs`, by running the following batch file:
`install_path\NetBackup\Logs\mklogdir.bat`.

Media servers have only the `bpbrm`, `bpcd`, `bpdm`, and `bptm` debug logs.

Creating NetBackup log file directories

Before you configure a feature that uses the OpenStorage framework, 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 OpenStorage. 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.

See the *NetBackup Troubleshooting Guide*.

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 251.

See [“About unified logging”](#) on page 248.

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 Troubleshooting Guide](#).

More information about the VxUL log system is available.

See the [NetBackup Troubleshooting Guide](#).

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 11-2](#).

[Table 11-2](#) describes the log files for each component.

Table 11-2 Logs for NetBackup MSDP activity

Component	VxUL OID	Description
Backups and restores	117	The nbj _m 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"> ■ The bpbr_m backup and restore manager. The following is the path to the log files: UNIX: /usr/opensv/netbackup/logs/bpbr_m Windows: <i>install_path</i>\Veritas\NetBackup\logs\bpbr_m ■ The bpd_{bm} database manager. The following is the path to the log files: UNIX: /usr/opensv/netbackup/logs/bpd_{bm} Windows: <i>install_path</i>\Veritas\NetBackup\logs\bpdb_m ■ The bpt_m tape manager for I/O operations. The following is the path to the log files: UNIX: /usr/opensv/netbackup/logs/bpt_m Windows: <i>install_path</i>\Veritas\NetBackup\logs\bpt_m

Table 11-2 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"> ■ For <code>bptm</code>: UNIX: <code>/usr/opensv/netbackup/logs/bptm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bptm</code> ■ For <code>bpstsinfo</code>: Windows: <code>/usr/opensv/netbackup/logs/admin</code> UNIX: <code>/usr/opensv/netbackup/logs/bpstsinfo</code> Windows: <code>install_path\Veritas\NetBackup\logs\admin</code> Windows: <code>install_path\Veritas\NetBackup\logs\stsinfo</code> ■ For <code>bpbrm</code>: UNIX: <code>/usr/opensv/netbackup/logs/bpbrm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bpbrm</code>
Client deduplication proxy server	N/A	<p>The deduplication proxy server <code>nbostpxy</code> on the client writes messages to files in the following directory, as follows:</p> <p>UNIX: <code>/usr/opensv/netbackup/logs/nbostpxy</code></p> <p>Windows: <code>install_path\Veritas\NetBackup\logs\nbostpxy</code>.</p>
Deduplication configuration script	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"> ■ UNIX: <code>storage_path/log/pdde-config.log</code> ■ Windows: <code>storage_path\log\pdde-config.log</code> <p>NetBackup creates this log file during the configuration process. If your configuration succeeded, you do not need to examine the log file. The only reason to look at the log file is if the configuration failed. If the configuration process fails after it creates and populates the storage directory, this log file identifies when the configuration failed.</p>
Deduplication 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 “About the MSDP pd.conf configuration file” on page 156.</p> <p>See “Editing the MSDP pd.conf file” on page 156.</p>
Device configuration and monitoring	111	The <code>nbemm</code> process.

Table 11-2 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.
Installation	N/A	The NetBackup installation process writes information about the installation of the deduplication components to a log file in the following directory: <ul style="list-style-type: none"> ■ UNIX: /var/log/puredisk ■ Windows: %ALLUSERSPROFILE%\Symantec\NetBackup\InstallLogs
NetBackup Deduplication Engine	N/A	The NetBackup Deduplication Engine writes several log files, as follows: <ul style="list-style-type: none"> ■ Log files in the <i>storage_path/log/spoold</i> directory, as follows: <ul style="list-style-type: none"> ■ The <i>spoold.log</i> file is the main log file ■ The <i>storaged.log</i> file is for queue processing. ■ 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.symantec.org/crcontrol/Control/010112.log</i> Usually, the only reason to examine these connection log files is if a Symantec support representative asks you to. ■ A VxUL log file for the events and errors that NetBackup receives from polling. The originator ID for the deduplication engine is 364.
NetBackup Deduplication Engine	364	The NetBackup Deduplication Engine that runs on the deduplication storage server.

Table 11-2 Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
NetBackup Deduplication Manager	N/A	<p>The log files are in the <code>/storage_path/log/spad</code> directory, as follows:</p> <ul style="list-style-type: none"> ■ <code>spad.log</code> ■ <code>sched_QueueProcess.log</code> ■ <code>SchedClass.log</code> ■ A log file for each connection to the manager is stored in a directory in the storage path <code>spad</code> directory. The following describes the pathname to a log file for a connection: <code>hostname/application/TaskName/MMDDYY.log</code> For example, the following is an example of a <code>bpstsinfo</code> connection log pathname on a Linux system: <code>/storage_path/log/spoold/server.symantecs.org/bpstsinfo/spad/010112.log</code> <p>Usually, the only reason to examine these connection log files is if a Symantec support representative asks you to.</p> <p>You can set the log level and retention period in the Change Storage Server dialog box Properties tab.</p> <p>See “Changing MSDP storage server properties” on page 196.</p>
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"> ■ The NetBackup <code>bptm</code> 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> ■ The following is the path name of MSDP replication log file: <code>/storage_path/log/spad/replication.log</code>
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>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 OID 387 to 2 or less. To configure unified logs, see the following guide:</p> <p>The NetBackup Troubleshooting Guide.</p>

Table 11-2 Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
Resilient network connections	N/A	<p>The deduplication plug-in logs information about keeping the connection alive.</p> <p>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 “About the MSDP pd.conf configuration file” on page 156.</p> <p>See “Editing the MSDP pd.conf file” on page 156.</p>

Troubleshooting MSDP installation issues

The following sections may help you troubleshoot installation issues.

See [“MSDP installation on SUSE Linux fails”](#) on page 258.

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 253.

See [“MSDP storage server configuration fails”](#) on page 259.

See [“MSDP database system error \(220\)”](#) on page 259.

See [“MSDP server not found error”](#) on page 260.

See [“License information failure during MSDP configuration”](#) on page 260.

See [“The disk pool wizard does not display an MSDP volume”](#) on page 261.

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 175.

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 253..

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: getStoreageServerInfo`

`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 262.

- See [“MSDP backup or duplication job fails”](#) on page 262.
- See [“MSDP client deduplication fails”](#) on page 264.
- See [“MSDP volume state changes to DOWN when volume is unmounted”](#) on page 265.
- See [“MSDP errors, delayed response, hangs”](#) on page 266.
- See [“Cannot delete an MSDP disk pool”](#) on page 266.
- See [“MSDP media open error \(83\)”](#) on page 267.
- See [“MSDP media write error \(84\)”](#) on page 269.
- See [“MSDP no images successfully processed \(191\)”](#) on page 272.
- See [“MSDP storage full conditions”](#) on page 272.
- See [“Specifying the restore server”](#) on page 227.

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 26.

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 273.

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 216.

See [“About MSDP log files”](#) on page 253.

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 253.

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 38.

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       : @aaaab
Total Capacity (GB) : 49.98
Free Space (GB)     : 43.66
Use%                : 12
Status              : UP
Flag                : ReadOnWrite
Flag                : AdminUp
Flag                : InternalUp
Num Read Mounts     : 0
Num Write Mounts    : 1
Cur Read Streams   : 0
Cur Write Streams  : 0
```

To change the volume state to UP

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 213.

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 26.

For virtual machines, Symantec recommends that you do the following:

- Set the memory size of each virtual machine to double the physical memory of the host.
- Set the minimum and the maximum values of each virtual machine to the same value (double the physical memory of the host). These memory settings prevent the virtual memory from becoming fragmented on the disk because it does not grow or shrink.

These recommendations may not be the best configuration for every virtual machine. However, Symantec recommends that you try this solution first when troubleshooting performance issues.

Cannot delete 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: `install_path\NetBackup\bin\admincmd\nbstlutil stlilist -incomplete`

Windows: `/usr/opensv/netbackup/bin/admincmd/nbstlutil stlilist -incomplete`

- 2 Cancel the incomplete jobs by running the following command for each backup ID returned by the previous command (xxxxx represents the backup ID):

UNIX: `install_path\NetBackup\bin\admincmd\nbstlutil cancel -backupid xxxxx`

Windows: `/usr/opensv/netbackup/bin/admincmd/nbstlutil cancel -backupid xxxxx`

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 36.

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 client-side backups fail

Client-side backups of an SQL 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 253.

Restore or duplication fails

The `media open error (83)` message appears in the **NetBackup Administration Console Activity Monitor**.

[Table 11-3](#) describes other messages that may appear.

Table 11-3 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
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 (PureDisk source)	Image open failed: error 2060018: file not found	<code>sts_open_image</code> failed: error 2060018
Replication (PureDisk 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 11-4](#) 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](#)
- [Duplication job to PureDisk storage fails](#)

Table 11-4 Media write error causes

The NetBackup Deduplication Engine (spoold) was too busy to respond.	Examine the Disk Logs report for PureDisk errors. Examine the disk monitoring services log files for details from the deduplication plug-in. See “Viewing MSDP disk reports” on page 189.
Data removal is running.	Data cannot be backed up at the same time as it is removed. See “About MSDP queue processing” on page 216.
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 “About adding additional MSDP storage” on page 51.
The deduplication pool is down.	Change the state to up. See “Changing a Media Server Deduplication Pool state” on page 205.
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 175.

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 38.

Duplication job to PureDisk storage fails

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.symantecs.org) 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.symantecs.org backup id msdp.symantecs.org_1360303271 optimized
  duplication failed, media write error (84).
2/12/2013 2:33:39 AM - Info msdp.symantecs.org(pid=25204)
  StorageServer=PureDisk:msdp.symantecs.org; Report=PDDO Stats for
  (msdp.symantecs.org): 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.symantecs.org_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 11-5](#) describes other messages that may appear.

Table 11-5 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	<code>sts_open_image</code> 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 185.

See [“About MSDP container files”](#) on page 187.

See [“Viewing storage usage within MSDP container files”](#) on page 187.

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 216.

See [“About the MSDP data removal process”](#) on page 225.

Although not advised, you can reclaim free space manually.

See [“Processing the MSDP transaction queue manually”](#) on page 216.

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 189.
- The NetBackup `bpererror` command with the `-disk` option reports on disk errors. The command resides in the following directories:
 UNIX: `/usr/opensv/netbackup/bin/admincmd`
 Windows: `install_path\Veritas\NetBackup\bin\admincmd`

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 11-6 MSDP event codes and messages

Event #	Event Severity	NetBackup Severity	Message example
1000	2	Error	Operation configload/reload failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1001	2	Error	Operation configload/reload failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1002	4	Warning	The open file limit exceeded in server PureDisk:server1.symantecs.org on host server1.symantecs.org. Will attempt to continue further.
1003	2	Error	A connection request was denied on the server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1004	1	Critical	Network failure occurred in server PureDisk:server1.symantecs.org on host server1.symantecs.org.

Table 11-6 MSDP event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1013	1	Critical	Task session start request on server PureDisk:server1.symantecs.org on host server1.symantecs.org got an unexpected error.
1008	2	Error	Task Aborted; An unexpected error occurred during communication with remote system in server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1009	8	Authorization	Authorization request from <IP> for user <USER> denied (<REASON>).
1010	2	Error	Task initialization on server PureDisk:server1.symantecs.org on host server1.symantecs.org got an unexpected error.
1011	2	Error	Task ended on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1012	2	Error	A request for agent task was denied on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1014	1	Critical	Task session start request on server PureDisk:server1.symantecs.org on host server1.symantecs.org got an unexpected error.
1015	1	Critical	Task creation failed, could not initialize task class on server PureDisk:server1.symantecs.org on host server1.symantecs.org.

Table 11-6 MSDP event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1017	1	Critical	Service Symantec DeduplicationEngine exit on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Please check the server log for the probable cause of this error. The application has terminated.
1018	16	Info	Startup of Symantec Deduplication Engine completed successfully on server1.symantecs.org.
1019	1	Critical	Service Symantec DeduplicationEngine restart on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Please check the server log for the probable cause of this error. The application has restarted.
1020	1	Critical	Service Symantec Deduplication Engine connection manager restart failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Please check the server log for the probable cause of this error.The application has failed to restart.
1028	1	Critical	Service Symantec DeduplicationEngine abort on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Please check the server log for the probable cause of this error.The application has caught an unexpected signal.
1029	1	Critical	Double backend initialization failure; Could not initialize storage backend or cache failure detected on host PureDisk:server1.symantecs.org in server server1.symantecs.org.

Table 11-6 MSDP event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1030	1	Critical	Operation Storage Database Initialization failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1031	1	Critical	Operation Content router context initialization failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1032	1	Critical	Operation log path creation/print failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1036	4	Warning	Operation a transaction failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1037	4	Warning	Transaction failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org. Transaction will be retried.
1040	2	Error	Operation Database recovery failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1043	2	Error	Operation Storage recovery failed on server PureDisk:server1.symantecs.org on host server1.symantecs.org.
1044	multiple	multiple	The usage of one or more system resources has exceeded a warning level. Operations will or could be suspended. Please take action immediately to remedy this situation.

Table 11-6 MSDP event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1057			<p>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.</p> <p>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.</p>
2000		Error	<p>Low space threshold exceeded on the partition containing the storage database on server PureDisk:server1.symantecs.org on host server1.symantecs.org.</p>

NetBackup appliance deduplication

This appendix includes the following topics:

- [About NetBackup appliance deduplication](#)
- [About Fibre Channel to a NetBackup 5020 appliance](#)
- [Enabling Fibre Channel to a NetBackup 5020 appliance](#)
- [Disabling Fibre Channel to a NetBackup 5020 appliance](#)
- [Displaying NetBackup 5020 appliance Fibre Channel port information](#)

About NetBackup appliance deduplication

NetBackup appliances are hardware and software solutions from Symantec that combine a host and storage with the Symantec backup software. The appliances offer customers easy and convenient deployment options for Symantec's industry-leading backup and deduplication technologies. The appliances enable efficient, storage-optimized data protection for the datacenter, remote office, and virtual environments.

Symantec's NetBackup appliance family consists of the following two series:

- NetBackup 5200 series of enterprise backup appliances that is based on the NetBackup backup platform. The 5200 series provides between 4 TB to 32 TB of deduplication storage.
A NetBackup 5200 series appliance can be a destination for optimized duplication from a NetBackup **Media Server Deduplication Pool**.

- NetBackup 5000 series of scalable deduplication appliances that is based on the PureDisk backup platform. The 5000 series is scalable from 16 TB to 192 TB of storage.
 A NetBackup 5000 series can be a storage destination for both the NetBackup Client Deduplication Option and the NetBackup Media Server Deduplication Option.

The NetBackup appliances share many common features, as follows:

- Easy to install, configure, and use.
- Modular capacity to fulfill your storage needs.
- A solution for the datacenter, remote office and branch office, and virtual machine backups. Source or target deduplication. Optimized synthetic backup to minimize data movement. Tape support for long-term data retention.
- Built in disk to disk replication for disaster recovery and an alternative solution to tape based vaulting
- Enterprise-class hardware and software. Hardware monitoring with a call home feature.

About Fibre Channel to a NetBackup 5020 appliance

With this release of NetBackup, you can use Fibre Channel for NetBackup MSDP storage server data traffic to a Symantec NetBackup 5020 appliance. The appliance must be at software release 1.3 or later.

The 5020 appliance provides the same functionality as a traditional PureDisk environment, which is a software only environment that is installed on your hardware. This appliance solution can function as the PureDisk storage destination for the following operations:

- A target for NetBackup client backups.
- A target for optimized duplication from a NetBackup **Media Server Deduplication Pool**.
- A source for optimized duplication to another deduplication destination.

For any operation that involves NetBackup PureDisk in this guide, PureDisk means both a traditional PureDisk environment and a NetBackup 5020 appliance.

Requirements for the NetBackup MSDP storage server:

- NetBackup 7.5 or later.
- A supported operating system.

See the NetBackup operating system compatibility list on the [NetBackup landing page](#) on the Symantec support Web site.

- One Qlogic 2562 (ISP 2532) HBA with 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 182.

For information about configuring the appliance and zoning the Fibre Channel SAN, see the appliance documentation.

Enabling Fibre Channel to a NetBackup 5020 appliance

You can enable Fibre Channel 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 279.

To enable 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/opensv/pdde/pdconfigure/scripts/support/dedup_fcmanager.sh -e
```

```
WARNING: Enabling/disabling Fibre Channel transport may require  
spad to be restarted.
```

```
Do you want to continue? [y/n]
```

- 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 157.

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 279.

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 157.

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 283.
- Statistics.
See [“To display Fibre Channel statistics on a NetBackup 5000 series appliance”](#) on page 284.

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 279.

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.symantecs.org)
/dev/sg5    SYMANTEC    192.168.0.3(5020-Silver.symantecs.org)

**** VLAN ****
The result is based on the scan at Sun, Jan 1 00:00:01 CST 2012
/dev/sg3    192.168.0.2
/dev/sg8    192.168.1.2
/dev/sg5    192.168.0.3

**** Fibre Channel Transport ****
Replication over Fibre Channel is disabled
Backup/Restore over Fibre Channel is disabled
```

This output shows the target mode Fibre Channel ports and the hosts to which Fibre Channel traffic can travel.

To display Fibre Channel statistics on a NetBackup 5000 series appliance

- ◆ As the root user, run the `dedup_fcmanager.sh` script with the `-t` option and interval and repeat arguments. The following command example lists the statistics five times with a one second interval between them:

```
usr/openv/pdde/pdconfigure/scripts/support/dedup_fcmanager.sh -t 1 5
```

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	0	0	0	0

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	2823	12702	0	17144

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	2105	9557	0	13070

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	2130	9597	0	13161

Port	I/O R(count/s)	I/O W(count/s)	I/O R(KB/s)	I/O W(KB/s)
5	0	0	0	0
6	2108	9632	0	13136

Some versions of some `qla2xxxx` drivers do not provide KB/s output.

Index

A

- about NetBackup deduplication 14
- about NetBackup deduplication options 14
- about the deduplication host configuration file 175
- Active secondary operations 145
- appliance deduplication 15–16
- attributes
 - clearing deduplication pool 212
 - clearing deduplication storage server 197
 - OptimizedImage 35
 - setting deduplication pool 206
 - setting deduplication storage server 195
 - viewing deduplication pool 206
 - viewing deduplication storage server 194
- Auto Image Replication
 - and trusted master servers 130
 - Backup operation in source domain 122–123
 - configuring MSDP replication to a different domain 120
 - for MSDP 132
 - overview 122
 - synchronizing clocks of master servers 123

B

- Backup
 - operation 144
- backup
 - client deduplication process 243
- big endian 230, 233
- bpstsinfo command 128
- byte order 230, 233

C

- cache hits field of the job details 184
- capacity
 - adding storage 51
- capacity and usage reporting for deduplication 185
- case sensitivity
 - in NetBackup names 21
- changing deduplication server hostname 198

- changing the deduplication storage server name and path 198
- CIFS 47
- clearing deduplication pool attributes 212
- client deduplication
 - about 28
 - components 243
 - disabling for a specific client 149
 - host requirements 30
- Common Internet File System 47
- compacting container files 187
- compression
 - and deduplication 33
 - pd.conf file setting 158
- configuring a deduplication pool 84
- configuring a deduplication storage server 72, 79
- configuring a deduplication storage unit 93
- configuring deduplication 56, 59
- container files
 - about 187
 - compaction 187
 - viewing capacity within 187
- contentrouter.cfg file
 - about 170
 - parameters for data integrity checking 221
 - ServerOptions for encryption 98
- Coordinated Universal Time (UTC) 123
- CR sent field of the job details 184
- credentials 31
 - adding NetBackup Deduplication Engine 203
 - changing NetBackup Deduplication Engine 204

D

- data classifications
 - in storage lifecycle policies 137, 144
 - use of Any 138, 191
- data integrity checking
 - about deduplication 217
 - configuration settings for deduplication 220
 - configuring behavior for deduplication 218

- data removal process
 - for deduplication 225
- database system error 259
- deactivating media server deduplication 235
- dedup field of the job details 184
- deduplication
 - about credentials 31
 - about fingerprinting 67
 - about the license key 54
 - adding credentials 203
 - and Fibre Channel 48
 - and iSCSI 48
 - cache hits field of the job details 184
 - capacity and usage reporting 185
 - changing credentials 204
 - client backup process 243
 - compression 33
 - configuration file 156
 - configuring 56, 59
 - configuring optimized synthetic backups 99
 - container files 187
 - CR sent field of the job details 184
 - data removal process 225
 - dedup field of the job details 184
 - encryption 34
 - event codes 273
 - license key for 54
 - licensing 54
 - limitations 28
 - media server process 240
 - network interface 32
 - node 26
 - performance 36
 - planning deployment 20
 - requirements for optimized within the same domain 104
 - scaling 38
 - scanned field of the job details 185
 - storage capacity 50
 - storage destination 22
 - storage management 51
 - storage paths 50
 - storage requirements 46
 - storage unit properties 95
- deduplication configuration file
 - editing 66, 156
 - parameters 157
- deduplication data integrity checking
 - about 217
 - deduplication data integrity checking *(continued)*
 - configuring behavior for 218
 - settings 220
- deduplication deduplication pool. *See* deduplication pool
- deduplication disk volume
 - changing the state 213
 - determining the state 213
- deduplication encryption
 - enabling 97
- deduplication host configuration file
 - about 175
 - deleting 175
- deduplication hosts
 - and firewalls 33
 - client requirements 30
 - load balancing server 26
 - server requirements 26
 - storage server 25
- deduplication installation
 - log file 256
- deduplication logs
 - about 253
 - client deduplication proxy plug-in log 255
 - client deduplication proxy server log 255
 - configuration script 255
 - deduplication plug-in log 255
 - NetBackup Deduplication Engine 256
 - NetBackup Deduplication Manager 257
 - VxUL deduplication logs 253
- deduplication node
 - about 26
 - adding a load balancing server 155
 - removing a load balancing server 199
- deduplication optimized synthetic backups
 - about 35
- deduplication plug-in
 - about 238
 - log file 255
- deduplication plug-in configuration file
 - configuring 61
- deduplication pool. *See* deduplication pool
 - about 83
 - changing properties 207
 - changing the state 205
 - clearing attributes 212
 - configuring 84
 - deleting 215
 - determining the state 205

- deduplication pool (*continued*)
 - properties 91
 - setting attributes 206
 - viewing 205
 - viewing attributes 206
 - deduplication port usage
 - about 33
 - troubleshooting 270
 - deduplication processes do not start 262
 - deduplication rate
 - how file size affects 37
 - monitoring 181
 - deduplication reference database
 - about 238
 - deduplication registry
 - resetting 175
 - deduplication servers
 - about 25
 - components 237
 - host requirements 26
 - deduplication storage capacity
 - about 50
 - viewing capacity in container files 187
 - deduplication storage destination 22
 - deduplication storage paths 50
 - deduplication storage requirements 46
 - deduplication storage server
 - about 25
 - change the name 198
 - changing properties 196
 - clearing attributes 197
 - components 237
 - configuration failure 260
 - configuring 72, 79
 - defining target for Auto Image Replication 127
 - deleting 201
 - deleting the configuration 201
 - determining the state 193
 - editing configuration file 173
 - getting the configuration 172
 - recovery 233
 - replacing the host 229
 - setting attributes 195
 - setting the configuration 174
 - viewing 193
 - viewing attributes 194
 - deduplication storage server configuration file
 - about 171
 - deduplication storage server name
 - changing 198
 - deduplication storage type 22
 - Deduplication storage unit
 - Only use the following media servers 95
 - Use any available media server 95
 - deleting backup images 215
 - deleting deduplication host configuration file 175
 - disaster recovery
 - protecting the data 42
 - recovering the storage server after catalog recovery 235
 - disk failure
 - deduplication storage server 231
 - disk logs 189
 - disk logs report 187
 - disk pool
 - cannot delete 266
 - disk pool status report 186, 189
 - disk storage unit report 189
 - Disk type 95
 - disk volume
 - changing the state 213
 - determining the state of a deduplication 213
 - volume state changes to down 265
 - domains
 - replicating backups to another. See Auto Image Replication
 - Duplication
 - optimized 145
 - Duplication to remote master. See Auto Image Replication
- ## E
- Enable file recovery from VM backup 146
 - encryption
 - and deduplication 34
 - enabling for deduplication 97
 - pd.conf file setting 160
 - endian
 - big 230, 233
 - little 230, 233
 - event codes
 - deduplication 273
- ## F
- Fibre Channel
 - and iSCSI comparison 48

- Fibre Channel (*continued*)
 - support for 47
- Fibre Channel and iSCSI comparison for MSDP 48
- file system
 - CIFS 47
 - NFS 47
 - Veritas File System for deduplication storage 51
- fingerprinting
 - about deduplication 67
- firewalls and deduplication hosts 33
- FlashBackup policy
 - Maximum fragment size (storage unit setting) 95
- FQDN or IP Address property in Resilient Network
 - host properties 151

G

garbage collection. *See* queue processing

H

host requirements 26

I

- images on disk report 189
- Import
 - operation 139
- iSCSI
 - and Fibre Channel comparison 48
 - support for 47

J

job ID search in unified logs 251

L

- legacy logging 251
 - directories 252
 - locations 251
- license information failure
 - for deduplication 260
- license key
 - for deduplication 54
- licensing deduplication 54
- limitations
 - media server deduplication 28
- little endian 230, 233
- load balancing server
 - about 26
 - adding to a deduplication node 155

- load balancing server (*continued*)
 - for deduplication 26
 - removing from deduplication node 199
- logging
 - see* legacy logging 251
- logs
 - about deduplication 253
 - Auto Image Replication 257
 - client deduplication proxy plug-in log 255
 - client deduplication proxy server log 255
 - deduplication configuration script log 255
 - deduplication installation 256
 - deduplication plug-in log 255
 - disk 189
 - NetBackup Deduplication Engine log 256
 - NetBackup Deduplication Manager log 257
 - optimized duplication 257
 - VxUL deduplication logs 253

M

- maintenance processing. *See* queue processing
- Maximum concurrent jobs 96
- Maximum fragment size 95
- media server deduplication
 - process 240
- Media Server Deduplication Option
 - about 23
- Media Server Deduplication Pool 132
- media server deduplication pool. *See* deduplication pool
- migrating from PureDisk to NetBackup
 - deduplication 178
- migrating to NetBackup deduplication 179
- mklogdir.bat 252
- MSDP
 - replication 122
 - replication target, configuring 132
- MSDP replication
 - about 36
- MSDP storage rebasing. *See* rebasing
- mtstrm.conf file
 - configuring 61

N

- nbstserv process 140
- NDMP
 - storage units 145

- NetBackup
 - naming conventions 21
- NetBackup 5000 series appliance
 - as a storage destination 22
- NetBackup 5200 series appliance
 - as a storage destination 22
- NetBackup appliance deduplication 15
- NetBackup Client Deduplication Option 15
- NetBackup deduplication
 - about 14
 - license key for 54
- NetBackup Deduplication Engine
 - about 238
 - about credentials 31
 - adding credentials 203
 - changing credentials 204
 - logs 256
- NetBackup Deduplication Manager
 - about 238
 - logs 257
- NetBackup deduplication options 14
- NetBackup Media Server Deduplication Option 15
- network interface
 - for deduplication 32
- NFS 47
- node
 - deduplication 26

O

- OpenStorage
 - optimized duplication 145
- OpenStorage appliance deduplication 16
- OpenStorage Disk Option 137
- optimized deduplication
 - configuring bandwidth 119
 - configuring for MSDP 99
 - limitations 104
 - logs 257
 - pull configuration within the same domain 113
 - push configuration 109
 - separate network for 115
- optimized deduplication copy
 - guidance for 106
- optimized duplication
 - about 36
 - about the media server in common within the same domain 106
 - push configuration within the same domain 106

- optimized MSDP deduplication
 - requirements 104
 - within the same domain 103
- optimized synthetic backups
 - configuring for deduplication 99
 - deduplication 35
- OptimizedImage attribute 35

P

- pd.conf file
 - about 156
 - editing 66, 156
 - parameters 157
- pdde-config.log 255
- performance
 - deduplication 36
 - monitoring deduplication rate 181
- policies
 - changing properties 146
 - creating 146
- port usage
 - and deduplication 33
 - troubleshooting 270
- Postponed secondary operations 145
- Priority for secondary operations 144
- provisioning the deduplication storage 45
- PureDisk appliance deduplication 16
- PureDisk deduplication 15
- PureDisk Deduplication Option
 - replacing with media server deduplication 177

Q

- queue processing 216
 - invoke manually 216

R

- rebasing
 - about 223
 - FP_CACHE_PERIOD_REBASING_THRESHOLD
 - parameter 163
 - FP_CACHE_REBASING_THRESHOLD
 - parameter 164
 - RebaseMaxPercentage parameter 225
 - RebaseMaxTime parameter 225
 - RebaseMinContainers parameter 225
 - RebaseScatterThreshold parameter 225
 - server-side rebasing parameters 224

- recovery
 - deduplication storage server 233
 - from deduplication storage server disk failure 231
 - Red Hat Linux
 - deduplication processes do not start 262
 - replacing PDDO with NetBackup deduplication 177
 - replacing the deduplication storage server 229
 - replication
 - between NetBackup domains. *See* Auto Image Replication
 - configuring MSDP replication to a different domain 120
 - for MSDP 36, 122
 - to an alternate NetBackup domain. *See* Auto Image Replication
 - Replication Director
 - Policy Configuration Wizard, unsupported 146
 - reports
 - disk logs 187, 189
 - disk pool status 186, 189
 - disk storage unit 189
 - reqlib directory 252
 - resetting the deduplication registry 175
 - Resiliency property in Resilient Network host properties 151
 - Resilient connection
 - Resilient Network host properties 149
 - resilient network connection
 - log file 257
 - Resilient Network host properties 149
 - FQDN or IP Address property in 151
 - Resiliency property in 151
 - restores
 - at a remote site 226
 - how deduplication restores work 43
 - specifying the restore server 227
 - reverse host name lookup
 - prohibiting 260
 - reverse name lookup 260
- S**
- scaling deduplication 38
 - scanned field of the job details 185
 - server
 - NetBackup debug logs 252
 - server not found error 260
 - setting deduplication pool attributes 206
 - SLP Parameters 118
 - snapshots
 - operation type 144
 - spa.cfg file
 - parameters for data integrity checking 221
 - Status Collection Daemon 252
 - storage capacity
 - about 50
 - adding 51
 - for deduplication 50
 - viewing capacity in container files 187
 - Storage Lifecycle Manager service (nbstserv) 140
 - storage lifecycle policies
 - Active secondary operations 145
 - best practices document 138
 - Data classification setting 144
 - hierarchy 142
 - operations 140
 - Postponed secondary operations 145
 - Priority for secondary operations 144
 - Storage lifecycle policy name 144
 - utility 137
 - Validate Across Backup Policies button 145
 - storage paths
 - about reconfiguring 198
 - changing 198
 - for deduplication 50
 - storage rebasing. *See* rebasing
 - storage requirements
 - for deduplication 46
 - storage server
 - about the configuration file 171
 - change the name 198
 - changing properties for deduplication 196
 - changing the name 198
 - components for deduplication 237
 - configuring for deduplication 72, 79
 - deduplication 25
 - define target for Auto Image Replication 127
 - deleting a deduplication 201
 - deleting the deduplication configuration 201
 - determining the state of a deduplication 193
 - editing deduplication configuration file 173
 - getting deduplication configuration 172
 - recovery 233
 - replacing the deduplication host 229
 - setting the deduplication configuration 174
 - viewing 193
 - viewing attributes 194

- storage server configuration
 - getting 172
 - setting 174
- storage server configuration file
 - editing 173
- storage type
 - for deduplication 22
- storage unit
 - configuring for deduplication 93
 - properties for deduplication 95
 - recommendations for deduplication 96
- storage unit groups
 - and storage lifecycle polices 145
 - not supported for Auto Image Replication source 122–123
- Storage unit name 95
- Storage unit type 95
- stream handlers
 - NetBackup 37
- synthetic backups
 - no Auto Image Replication support 123

T

- topology of storage 127–128
- troubleshooting
 - database system error 259
 - deduplication backup jobs fail 262
 - deduplication processes do not start 262
 - general operational problems 266
 - host name lookup 260
 - installation fails on Linux 258
 - no volume appears in disk pool wizard 261
 - server not found error 260
- trusted master servers
 - for Auto Image Replication 130

U

- unified logging 248
 - format of files 249
 - location 248
- uninstalling media server deduplication 235
- UTC (Coordinated Universal Time) 123

V

- Validation Report tab 145
- viewing deduplication pool attributes 206
- viewing storage server attributes 194
- VM backup 146

- vmscd 252
- volume manager
 - Veritas Volume Manager for deduplication storage 51
- vxlogview command 249
 - with job ID option 251

W

- wizards
 - Policy Configuration 146