

Veritas NetBackup™ Deduplication Guide

UNIX, Windows, and Linux

Release 8.3



Veritas NetBackup™ Deduplication Guide

Last updated: 2020-07-24

Legal Notice

Copyright © 2020 Veritas Technologies LLC. All rights reserved.

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

This product may contain third-party software for which Veritas 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. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC 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. Veritas Technologies LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

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

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introducing the NetBackup media server deduplication option	14
	About the NetBackup deduplication options	14
Chapter 2	Planning your deployment	16
	Planning your MSDP deployment	17
	NetBackup naming conventions	18
	About MSDP deduplication nodes	19
	About the NetBackup deduplication destinations	19
	About MSDP storage capacity	20
	About MSDP storage and connectivity requirements	21
	Fibre Channel and iSCSI comparison for MSDP	23
	About NetBackup media server deduplication	25
	About MSDP storage servers	27
	About MSDP load balancing servers	28
	About MSDP server requirements	28
	About MSDP unsupported configurations	30
	About NetBackup Client Direct deduplication	30
	About MSDP client deduplication requirements and limitations	32
	About MSDP remote office client deduplication	32
	About MSDP remote client data security	33
	About remote client backup scheduling	33
	About the NetBackup Deduplication Engine credentials	34
	About the network interface for MSDP	35
	About MSDP port usage	35
	About MSDP optimized synthetic backups	36
	About MSDP and SAN Client	37
	About MSDP optimized duplication and replication	37
	About MSDP performance	38
	How file size may affect the MSDP deduplication rate	39
	About MSDP stream handlers	39
	Oracle stream handler	39
	Microsoft SQL Server stream handler	40
	MSDP deployment best practices	41

	Use fully qualified domain names	41
	About scaling MSDP	41
	Send initial full backups to the storage server	42
	Increase the number of MSDP jobs gradually	43
	Introduce MSDP load balancing servers gradually	43
	Implement MSDP client deduplication gradually	43
	Use MSDP compression and encryption	44
	About the optimal number of backup streams for MSDP	44
	About storage unit groups for MSDP	44
	About protecting the MSDP data	45
	Save the MSDP storage server configuration	46
	Plan for disk write caching	46
Chapter 3	Provisioning the storage	47
	About provisioning the storage for MSDP	47
	Do not modify MSDP storage directories and files	49
	About adding additional MSDP storage	49
	About volume management for NetBackup MSDP	50
Chapter 4	Licensing deduplication	51
	About the MSDP license	51
	Licensing NetBackup MSDP	52
Chapter 5	Configuring deduplication	53
	Configuring MSDP server-side deduplication	55
	Configuring MSDP client-side deduplication	58
	About the MSDP Deduplication Multi-Threaded Agent	59
	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	68
	Configuring the MSDP fingerprint cache behavior	69
	MSDP fingerprint cache behavior options	69
	About seeding the MSDP fingerprint cache for remote client deduplication	70
	Configuring MSDP fingerprint cache seeding on the client	72
	Configuring MSDP fingerprint cache seeding on the storage server	73
	NetBackup seedutil options	75

Enabling 250-TB support for MSDP	76
About MSDP Encryption using NetBackup KMS service	76
Upgrading KMS for MSDP	77
About MSDP Encryption using external KMS server	79
Configuring a storage server for a Media Server Deduplication Pool	80
MSDP storage path properties	94
MSDP network interface properties	97
Configuring a storage server for a PureDisk Deduplication Pool	97
About disk pools for NetBackup deduplication	107
Configuring a disk pool for deduplication	109
Media Server Deduplication Pool properties	116
Creating the data directories for 250-TB MSDP support	118
Adding volumes to a 250-TB Media Server Deduplication Pool	119
Configuring a Media Server Deduplication Pool storage unit	121
Media Server Deduplication Pool storage unit properties	123
MSDP storage unit recommendations	124
Configuring client attributes for MSDP client-side deduplication	125
Disabling MSDP client-side deduplication for a client	127
About MSDP compression	127
About MSDP encryption	129
MSDP compression and encryption settings matrix	130
Configuring encryption for MSDP backups	132
Configuring encryption for MSDP optimized duplication and replication	133
About the rolling data conversion mechanism for MSDP	134
Modes of rolling data conversion	135
MSDP encryption behavior and compatibilities	137
Configuring optimized synthetic backups for MSDP	138
About a separate network path for MSDP duplication and replication	139
Configuring a separate network path for MSDP duplication and replication	140
About MSDP optimized duplication within the same domain	141
About the media servers for MSDP optimized duplication within the same domain	144
Configuring MSDP optimized duplication within the same NetBackup domain	149
Configuring NetBackup optimized duplication or replication behavior	153
About MSDP replication to a different domain	156
Configuring MSDP replication to a different NetBackup domain	158
About NetBackup Auto Image Replication	159

About trusted master servers for Auto Image Replication	167
About the certificate to be used for adding a trusted master server	171
Adding a trusted master server using a NetBackup CA-signed (host ID-based) certificate	173
Adding a trusted master server using external CA-signed certificate	178
Removing a trusted master server	179
Enabling NetBackup clustered master server inter-node authentication	180
Configuring NetBackup CA and NetBackup host ID-based certificate for secure communication between the source and the target MSDP storage servers	182
Configuring external CA for secure communication between the source MSDP storage server and the target MSDP storage server	183
Configuring a target for MSDP replication to a remote domain	184
About configuring MSDP optimized duplication and replication bandwidth	191
About storage lifecycle policies	192
About the storage lifecycle policies required for Auto Image Replication	192
Creating a storage lifecycle policy	194
Storage Lifecycle Policy dialog box settings	197
About MSDP backup policy configuration	199
Creating a backup policy	200
Resilient Network properties	200
Resilient connection resource usage	203
Specifying resilient connections	204
Adding an MSDP load balancing server	206
About variable-length deduplication on NetBackup clients	207
About the MSDP pd.conf configuration file	209
Editing the MSDP pd.conf file	210
MSDP pd.conf file parameters	210
About the MSDP contentrouter.cfg file	227
About saving the MSDP storage server configuration	227
Saving the MSDP storage server configuration	228
Editing an MSDP storage server configuration file	229
Setting the MSDP storage server configuration	230
About the MSDP host configuration file	231
Deleting an MSDP host configuration file	232
Resetting the MSDP registry	232

About protecting the MSDP catalog	233
About the MSDP shadow catalog	233
About the MSDP catalog backup policy	234
Changing the MSDP shadow catalog path	236
Changing the MSDP shadow catalog schedule	237
Changing the number of MSDP catalog shadow copies	238
Configuring an MSDP catalog backup	239
MSDP drcontrol options	240
Updating an MSDP catalog backup policy	244
About MSDP FIPS compliance	245
Configuring the NetBackup client-side deduplication to support multiple interfaces of MSDP	247
About MSDP multi-domain support	248
About MSDP multi-domain VLAN Support	251

Chapter 6

Configuring deduplication to the cloud with NetBackup Cloud Catalyst	253
Using NetBackup Cloud Catalyst to upload deduplicated data to the cloud	254
Optimized duplication is used to copy data from an MSDP storage server to a Cloud Catalyst storage server (preferred use case)	254
Backups go directly to a Cloud Catalyst storage server	256
Cloud Catalyst requirements and limitations	257
Configuring a Linux media server as a Cloud Catalyst storage server	259
Configuring a Cloud Catalyst storage server for deduplication to the cloud	261
How to configure a NetBackup Cloud Catalyst Appliance	261
How to configure a Linux media server as a Cloud Catalyst storage server	265
Configuring a Cloud Catalyst storage server as the target for the deduplications from MSDP storage servers	267
Certificate validation using Online Certificate Status Protocol (OCSP)	269
Managing Cloud Catalyst storage server with IAM Role or CREDS_CAPS credential broker type	270
Configuring a storage lifecycle policy for NetBackup Cloud Catalyst	270
About the Cloud Catalyst esfs.json configuration file	271
About the Cloud Catalyst cache	275
Controlling data traffic to the cloud when using Cloud Catalyst	276

	Configuring source control or target control optimized duplication for Cloud Catalyst	277
	Configuring a Cloud Catalyst storage server as the source for optimized duplication	283
	Decommissioning Cloud Catalyst cloud storage	283
	NetBackup Cloud Catalyst workflow processes	284
	Disaster recovery for Cloud Catalyst	289
	About image sharing in cloud using Cloud Catalyst	289
Chapter 7	MSDP cloud support	298
	About MSDP cloud support	298
	Operating system requirement for configuration	299
	Limitations	300
	Creating a cloud storage unit	300
	Updating cloud credentials for a cloud LSU	304
	Updating encryption configurations for a cloud LSU	305
	Deleting a cloud LSU	306
	Backup data to cloud by using cloud LSU	307
	Duplicate data cloud by using cloud LSU	308
	Configuring AIR to use cloud LSU	308
	About backward compatibility support	311
	About the configuration items in cloud.json, contentrouter.cfg and spa.cfg	312
	About the tool updates for cloud support	316
	About the disaster recovery for cloud LSU	318
	About Image Sharing using MSDP cloud	323
Chapter 8	Monitoring deduplication activity	325
	Monitoring the MSDP deduplication and compression rates	325
	Viewing MSDP job details	327
	MSDP job details	328
	About MSDP storage capacity and usage reporting	330
	About MSDP container files	332
	Viewing storage usage within MSDP container files	332
	Viewing MSDP disk reports	334
	About monitoring MSDP processes	335
	Reporting on Auto Image Replication jobs	335
Chapter 9	Managing deduplication	337
	Managing MSDP servers	338
	Viewing MSDP storage servers	338

Determining the MSDP storage server state	338
Viewing MSDP storage server attributes	339
Setting MSDP storage server attributes	340
Changing MSDP storage server properties	341
Clearing MSDP storage server attributes	342
About changing the MSDP storage server name or storage path	343
Changing the MSDP storage server name or storage path	343
Removing an MSDP load balancing server	345
Deleting an MSDP storage server	346
Deleting the MSDP storage server configuration	347
Managing NetBackup Deduplication Engine credentials	348
Determining which media servers have deduplication credentials	348
Adding NetBackup Deduplication Engine credentials	348
Changing NetBackup Deduplication Engine credentials	349
Deleting credentials from a load balancing server	349
Managing Media Server Deduplication Pools	350
Viewing Media Server Deduplication Pools	350
Determining the Media Server Deduplication Pool state	350
Changing Media Server Deduplication Pool state	351
Viewing Media Server Deduplication Pool attributes	351
Setting a Media Server Deduplication Pool attribute	352
Changing a Media Server Deduplication Pool properties	353
Clearing a Media Server Deduplication Pool attribute	358
Determining the MSDP disk volume state	359
Changing the MSDP disk volume state	359
Inventorying a NetBackup disk pool	360
Deleting a Media Server Deduplication Pool	361
Deleting backup images	361
About MSDP queue processing	362
Processing the MSDP transaction queue manually	362
About MSDP data integrity checking	363
Configuring MSDP data integrity checking behavior	364
MSDP data integrity checking configuration parameters	366
About managing MSDP storage read performance	368
About MSDP storage rebasing	368
MSDP server-side rebasing parameters	370
About the MSDP data removal process	371
Resizing the MSDP storage partition	372
How MSDP restores work	373
Configuring MSDP restores directly to a client	374
About restoring files at a remote site	374

	About restoring from a backup at a target master domain	375
	Specifying the restore server	375
Chapter 10	Recovering MSDP	377
	About recovering the MSDP catalog	377
	Restoring the MSDP catalog from a shadow copy	378
	Recovering from an MSDP storage server disk failure	380
	Recovering from an MSDP storage server failure	381
	Recovering the MSDP storage server after NetBackup catalog recovery	384
Chapter 11	Replacing MSDP hosts	385
	Replacing the MSDP storage server host computer	385
Chapter 12	Uninstalling MSDP	388
	About uninstalling MSDP	388
	Deactivating MSDP	388
Chapter 13	Deduplication architecture	390
	MSDP server components	390
	Media server deduplication backup process	393
	MSDP client components	396
	MSDP client-side deduplication backup process	396
Chapter 14	Troubleshooting	401
	About unified logging	401
	About using the <code>vxlogview</code> command to view unified logs	402
	Examples of using <code>vxlogview</code> to view unified logs	404
	About legacy logging	406
	Creating NetBackup log file directories for MSDP	406
	NetBackup MSDP log files	407
	Troubleshooting MSDP installation issues	413
	MSDP installation on SUSE Linux fails	413
	Troubleshooting MSDP configuration issues	414
	MSDP storage server configuration fails	414
	MSDP database system error (220)	414
	MSDP server not found error	415
	License information failure during MSDP configuration	416
	The disk pool wizard does not display an MSDP volume	416
	Troubleshooting MSDP operational issues	417

Verify that the MSDP server has sufficient memory	417
MSDP backup or duplication job fails	417
MSDP client deduplication fails	420
MSDP volume state changes to DOWN when volume is unmounted	420
MSDP errors, delayed response, hangs	421
Cannot delete an MSDP disk pool	422
MSDP media open error (83)	423
MSDP media write error (84)	425
MSDP no images successfully processed (191)	428
MSDP storage full conditions	428
Troubleshooting MSDP catalog backup	429
Storage Platform Web Service (spws) does not start	430
Disk volume API or command line option does not work	430
Viewing MSDP disk errors and events	430
MSDP event codes and messages	431
Troubleshooting Cloud Catalyst issues	433
Cloud Catalyst logs	434
Problems encountered while using the Cloud Storage Server Configuration Wizard	435
Disk pool problems	437
Problems during cloud storage server configuration	438
Status 191: No images were successfully processed	440
Media write error (84) if due to a full local cache directory	440
Troubleshooting restarting ESFS after the Cloud Catalyst storage server is down	441
Restarting the vxesfsd process	441
Problems restarting vxesfsd	443
Unable to create CloudCatalyst with a media server that has version earlier to 8.2	443
Cloud Catalyst troubleshooting tools	444
Unable to obtain the administrator password to use an AWS EC2 instance that has a Windows OS	446
Trouble shooting multi-domain issues	446
Unable to configure OpenStorage server from another domain	446
MSDP storage server is down when you configure an OpenStorage server	447
MSDP server is overloaded when it is used by multiple NetBackup domains	447

Appendix A	Migrating to MSDP storage	449
	Migrating from PureDisk to the NetBackup MSDP	449
	Migrating from another storage type to MSDP	450
Index		452

Introducing the NetBackup media server deduplication option

This chapter includes the following topics:

- [About the NetBackup deduplication options](#)

About the NetBackup deduplication options

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

Deduplication everywhere provides the following benefits:

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

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

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

Table 1-1 NetBackup deduplication options

Type	Description
Media server deduplication	<p>NetBackup clients send their backups to a NetBackup media server, which deduplicates the backup data. A NetBackup media server hosts the NetBackup Deduplication Engine, which writes the data to a Media Server Deduplication Pool on the target storage and manages the deduplicated data</p> <p>See “About NetBackup media server deduplication” on page 25.</p>
Client deduplication	<p>With NetBackup MSDP client deduplication, clients deduplicate their backup data and then send it directly to the storage server, which writes it to the storage. The network traffic is reduced greatly.</p> <p>See “About NetBackup Client Direct deduplication” on page 30.</p>
NetBackup appliance deduplication	<p>Veritas provides several hardware and a software solutions that include NetBackup deduplication.</p> <p>The NetBackup appliances have their own documentation set: https://www.veritas.com/content/support/en_US/Appliances.html</p>

Planning your deployment

This chapter includes the following topics:

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

Planning your MSDP deployment

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

Table 2-1 Deployment overview

Step	Deployment task	Where to find the information
Step 1	Learn about deduplication nodes and storage destinations	See “About MSDP deduplication nodes” on page 19. See “About the NetBackup deduplication destinations” on page 19.
Step 2	Understand the storage capacity and requirements	See “About MSDP storage capacity” on page 20. See “About MSDP storage and connectivity requirements” on page 21.
Step 3	Determine which type of deduplication to use	See “About NetBackup media server deduplication” on page 25. See “About NetBackup Client Direct deduplication” on page 30. See “About MSDP remote office client deduplication” on page 32.
Step 4	Determine the requirements for deduplication hosts	See “About MSDP storage servers” on page 27. See “About MSDP server requirements” on page 28. See “About MSDP client deduplication requirements and limitations” on page 32. See “About the network interface for MSDP” on page 35. See “About MSDP port usage” on page 35. See “About scaling MSDP” on page 41. See “About MSDP performance” on page 38.
Step 5	Determine the credentials for deduplication	See “About the NetBackup Deduplication Engine credentials” on page 34.
Step 6	Read about compression and encryption	See “About MSDP compression” on page 127. See “About MSDP encryption” on page 129.
Step 7	Read about optimized synthetic backups	See “About MSDP optimized synthetic backups” on page 36.
Step 8	Read about deduplication and SAN Client	See “About MSDP and SAN Client” on page 37.
Step 9	Read about optimized duplication and replication	See “About MSDP optimized duplication and replication” on page 37.
Step 10	Read about stream handlers	See “About MSDP stream handlers” on page 39.

Table 2-1 Deployment overview (*continued*)

Step	Deployment task	Where to find the information
Step 11	Read about best practices for implementation	See “MSDP deployment best practices” on page 41.
Step 12	Determine the storage requirements and provision the storage	See “About provisioning the storage for MSDP” on page 47. See “About MSDP storage and connectivity requirements” on page 21. See “About MSDP storage capacity” on page 20. See “MSDP storage path properties” on page 94.
Step 13	License MSDP	See “About the MSDP license” on page 51. See “Licensing NetBackup MSDP” on page 52.
Step 14	Configure MSDP	See “Configuring MSDP server-side deduplication” on page 55. See “Configuring MSDP client-side deduplication” on page 58.
Step 15	Migrate from other storage to NetBackup deduplication	See “Migrating from another storage type to MSDP” on page 450.

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 (_)

These characters are also used for foreign languages.

Note: No spaces are allowed.

The Logical Storage Unit (LSU) name or the Domain Volume name must have fewer than 50 ASCII characters including a hyphen (-) and an underscore (_) and must not have a blank space.

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

See [“About the NetBackup Deduplication Engine credentials”](#) on page 34.

About MSDP deduplication nodes

A media server deduplication node comprises the following:

Storage server	<p>The storage server deduplicates the backups, writes the data to the storage, and manages the storage.</p> <p>See “About MSDP storage servers” on page 27.</p>
Load balancing servers	<p>Load balancing servers assist the storage server by deduplicating backups. Load balancing servers are optional.</p> <p>See “About MSDP load balancing servers” on page 28.</p>
Storage	<p>See “About the NetBackup deduplication destinations” on page 19.</p>
Clients	<p>The clients may include the clients that deduplicate their own data (Client Direct).</p> <p>See “About NetBackup Client Direct deduplication” on page 30.</p>

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

Each node manages its own storage. Deduplication within each node is supported; deduplication between nodes is not supported.

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

See [“About MSDP storage servers”](#) on page 27.

About the NetBackup deduplication 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 or cloud 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>The Media Server Deduplication Pool can be hosted on the following systems:</p> <ul style="list-style-type: none"> ■ A NetBackup media server. ■ A NetBackup 5200 series appliance or NetBackup 5300 series appliance.
PureDisk Deduplication Pool	<p>A NetBackup PureDisk Deduplication Pool represents a storage pool on a NetBackup 5000 series appliance. NetBackup deduplicates the data, and the NetBackup 5000 series appliance hosts the storage.</p> <p>Note: NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p> <p>If you use a NetBackup 5000 series appliance for storage, use the following documentation:</p> <ul style="list-style-type: none"> ■ The NetBackup 5000 series appliance documentation to plan, implement, configure, and manage the appliance, which includes the storage: http://www.veritas.com/docs/TECH145661 ■ This guide to configure backups and deduplication in NetBackup. When you configure the storage server, select PureDisk Deduplication Pool as the storage type.

About MSDP storage capacity

The MSDP storage contains one local LSU or multiple cloud LSUs. The following table describes the maximum deduplication storage capacity for a single **Media Server Deduplication Pool** that contains only one local LSU:

Table 2-3 Maximum MSDP storage capacities

Maximum capacity	Description
64 TBs	For all supported systems, NetBackup supports up to 64 TBs of storage in a single Media Server Deduplication Pool . See “About provisioning the storage for MSDP” on page 47.
250 TBs	NetBackup supports 250 TBs of storage in a new Media Server Deduplication Pool on the supported versions of the following operating systems: <ul style="list-style-type: none"> ■ Red Hat Linux ■ Windows Server ■ SUSE Linux Recommended operating systems: <ul style="list-style-type: none"> ■ Red Hat Linux 7.5 ■ Windows Server 2012 R2 Datacenter See “About provisioning the storage for MSDP” on page 47.

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. If you use separate storage for the deduplication database, NetBackup still uses the 96-percent threshold to protect the data storage from any possible overload.

If your storage requirements exceed the capacity of a **Media Server Deduplication Pool**, you can use more than one media server deduplication node.

See [“About MSDP deduplication nodes”](#) on page 19.

For the operating system versions that NetBackup supports for deduplication, see the NetBackup operating system compatibility list available through the following URL:

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

About MSDP storage and connectivity requirements

The following subsections describe the storage and the connectivity requirements for the NetBackup Media Server Deduplication Option.

Storage media

The following are the minimum requirements for single stream read or write performance for each disk volume. Greater individual data stream capability or aggregate capability may be required to satisfy your objectives for writing to and reading from disk.

Up to 32 TBs of storage	130 MB/sec. 200 MB/sec for enterprise-level performance.
32 to 48 TBs of storage	200 MB/sec. Veritas recommends that you store the data and the deduplication database on separate disk volumes, each with 200 MB/sec read or write speed. Neither should be stored on the system disk.
48 to 64 TBs of storage	250 MB/sec. Veritas recommends that you store the data and the deduplication database on separate disk volumes, each with 250 MB/sec read or write speed. Neither should be stored on the system disk.
96 TBs of storage	250 MB/sec. 96 TBs of storage require four separate volumes, each with 250 MB/sec read or write speed. You cannot use the system disk of the storage server host for any of the required volumes.
250 TBs of storage	500 MB/sec

Local disk storage may leave you vulnerable in a disaster. SAN disk can be remounted at a newly provisioned server with the same name.

NetBackup requires the exclusive use of the disk resources. If the storage is also 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 the following storage types for deduplication storage:

- Network Attached Storage (that is, file based storage protocols) such as CIFS or NFS.
- The ZFS file system.

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

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

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

See [“About provisioning the storage for MSDP”](#) on page 47.

Storage connection

The storage must be direct-attached storage (DAS), internal disks, or connected by a dedicated, low latency storage area network (Fibre Channel or iSCSI).

A storage area network should conform to the following criteria:

Latency	Maximum 0.1-millisecond latency per round trip.
Bandwidth	<p>Enough bandwidth on the storage network to satisfy your throughput objectives.</p> <p>Veritas supports iSCSI on storage networks with at least 10-Gigabit Ethernet network bandwidth.</p> <p>Veritas recommends the Fibre Channel storage networks with at least 4-Gigabit network bandwidth.</p>
HBAs	<p>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>

See [“Fibre Channel and iSCSI comparison for MSDP”](#) on page 23.

Fibre Channel and iSCSI comparison for MSDP

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

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

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

Table 2-4 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, Veritas requires dedicated network interfaces on the storage server to reduce storage server load and reduce latency.
Latency	Low.	Higher.
Flow control	A built-in flow control mechanism that ensures data is sent to a device when it is ready to accept it.	No built-in flow control. Veritas recommends that you use the Ethernet priority-based flow control as defined in the IEEE 802.1Qbb standard.
Deployment	Difficult.	Easier than Fibre Channel, 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 Veritas supports iSCSI for connectivity to **Media Server Deduplication Pool** storage, Veritas recommends Fibre Channel. Veritas 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 423.


See [“MSDP media write error \(84\)”](#) on page 425.

About NetBackup media server deduplication

With media server deduplication, the NetBackup client software creates the image of backed up files as for a normal backup. The client sends the backup image to a media server, which hosts the plug-in that duplicates the backup data. The media server can be the storage server or a load balancing server if one is configured. The deduplication plug-in breaks the backup image into segments and compares the segments to all of the segments that are stored in that deduplication node. The plug-in then sends only the unique segments to the NetBackup Deduplication Engine on the storage server. The Deduplication Engine writes the data to a **Media Server Deduplication Pool**.

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

Figure 2-1 NetBackup media server deduplication



More detailed information is available.

See [“About MSDP deduplication nodes”](#) on page 19.

See [“About MSDP storage servers”](#) on page 27.

See [“About MSDP load balancing servers”](#) on page 28.

See [“About MSDP server requirements”](#) on page 28.

See [“About MSDP unsupported configurations”](#) on page 30.

See [“MSDP server components”](#) on page 390.

See [“Media server deduplication backup process”](#) on page 393.

About MSDP storage servers

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

See [“About MSDP deduplication nodes”](#) on page 19.

The MSDP storage server does the following:

- Receives the backups from clients and then deduplicates the data.
- Receives the deduplicated data from clients or from other media servers.
You can configure NetBackup clients and other NetBackup media servers to deduplicate data also. In which case, the storage server only receives the data after it is deduplicated.

See [“About NetBackup Client Direct deduplication”](#) on page 30.

See [“About MSDP load balancing servers”](#) on page 28.

- Writes the deduplicated data to and reads the deduplicated data from the disk or cloud storage.
- Manages that storage.
- Manages the deduplication processes.

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

- Optimized duplication between local LSUs 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 141.

- 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 159.

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 data segments 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 43.

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.

Note: Load balancing for Cloud Catalyst storage servers or Cloud Catalyst appliances is supported. See the following topic for configuration considerations.

See [“Configuring source control or target control optimized duplication for Cloud Catalyst”](#) on page 277.

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 2-5](#) 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 2-5 MSDP server minimum requirements

Component	Storage server	Load balancing server
CPU	Veritas recommends at least a 2.2-GHz clock rate. A 64-bit processor is required. At least four cores are required. Veritas recommends eight cores. For 64 TBs of storage, Intel x86-64 architecture requires eight cores.	Veritas recommends at least a 2.2-GHz clock rate. A 64-bit processor is required. At least two cores are required. Depending on throughput requirements, more cores may be helpful.
RAM	From 8 TBs to 32 TBs of storage, Veritas recommends 1GB RAM for 1TB of storage. However, beyond 32 TBs storage, Veritas recommends more than 32GBs of RAM for better and enhanced performance.	4 GBs.
Operating system	The operating system must be a supported 64-bit operating system. See the operating system compatibility list for your NetBackup release on the Veritas Support website. http://www.netbackup.com/compatibility	The operating system must be a supported 64-bit operating system. Load balancing for Cloud Catalyst storage servers or Cloud Catalyst appliances is supported. See the following topic for configuration considerations. See “Configuring source control or target control optimized duplication for Cloud Catalyst” on page 277. See the operating system compatibility list for your NetBackup release on the following website. http://www.netbackup.com/compatibility

A Veritas 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.

For more information, refer to <http://veritas.com/docs/TECH77575>.

Note: This page has been updated for NetBackup version 7.5.

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

See [“About MSDP performance”](#) on page 38.

See [“About MSDP queue processing”](#) on page 362.

About MSDP unsupported configurations

The following items describe some configurations that are not supported:

- NetBackup media server deduplication and Veritas Backup Exec deduplication cannot reside on the same host. If you use both NetBackup and Backup Exec deduplication, each product must reside on a separate host.
- NetBackup does not support clustering of deduplication storage servers or load balancing servers.
- Deduplication within each media server deduplication node is supported; global deduplication between nodes is not supported.

About NetBackup Client Direct deduplication

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

Client deduplication does the following:

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


NetBackup Client Deduplication is a solution for the following cases:

- Remote office or branch office backups to the data center.
 NetBackup provides resilient network connections for remote office backups.
 See [“About MSDP remote office client deduplication”](#) on page 32.
- LAN connected file server
- Virtual machine backups.

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

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

Figure 2-2 NetBackup client deduplication



More information is available.

See [“About MSDP client deduplication requirements and limitations”](#) on page 32.

See [“About MSDP remote office client deduplication”](#) on page 32.

See [“MSDP client components”](#) on page 396.

See [“MSDP client-side deduplication backup process”](#) on page 396.

About MSDP client deduplication requirements and limitations

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

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

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

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

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

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

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

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

About MSDP remote office client deduplication

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

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

- Client deduplication reduces the time that is required for WAN backups by reducing the amount of data that must be transferred.

- The resilient connections provide automatic recovery from network failures and latency (within the parameters from which NetBackup can recover).

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

See [“Resilient Network properties”](#) on page 200.

See [“Specifying resilient connections”](#) on page 204.

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 210.

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 70.

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. Veritas recommends that you use the deduplication encryption to protect your data during your remote client backups.

See [“About MSDP encryption”](#) on page 129.

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

See [“How MSDP restores work”](#) on page 373.

About remote client backup scheduling

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

About the NetBackup Deduplication Engine credentials

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

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

The following are the rules for the credentials:

- The user name and the password can be up to 62 characters in length. The user name and the password cannot be empty.
- 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.
 - Caret sign (^).
 - Percent sign (%).
 - United States dollar sign (\$).
 - Spaces.
 - Leading and trailing 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 Veritas support representative.

About the network interface for MSDP

If the MSDP storage server has more than one network interface, NetBackup uses the default interface for all deduplication traffic. (Deduplication traffic includes backups, restores, duplication, and replication.) The host operating system determines which network interface is the default. However, you can configure the network interface or interfaces that NetBackup uses, as follows:

Configure a specific interface	<p>To use a specific interface, you can enter that interface name when you configure the deduplication storage server. NetBackup uses this interface for all deduplication traffic unless you also configure a separate interface for duplication and replication.</p> <p>See “MSDP network interface properties” on page 97.</p> <p>See “Configuring a storage server for a Media Server Deduplication Pool” on page 80.</p>
Configure an interface for duplication and replication traffic	<p>You can configure a separate network interface for the duplication and the replication traffic. The backup and restore traffic continues to use the default interface or the specific configured interface.</p> <p>See “About a separate network path for MSDP duplication and replication” on page 139.</p> <p>See “Configuring a separate network path for MSDP duplication and replication” on page 140.</p>

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 2-6 Deduplication ports

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

About MSDP optimized synthetic backups

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

The optimized synthetic backup method provides the following benefits:

- Faster than a synthetic backup.
Regular synthetic backups are constructed on the media server. They are moved across the network from the storage server to the media server and synthesized into one image. The synthetic image is then moved back to the storage server.
- Requires no data movement across the network.
Regular synthetic backups use network traffic.

See [“Configuring optimized synthetic backups for MSDP”](#) on page 138.

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 340.

See [“Setting a Media Server Deduplication Pool attribute”](#) on page 352.

Table 2-7 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 2-8 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 141. See “About MSDP cloud support” on page 298.
To a remote NetBackup domain	See “About NetBackup Auto Image Replication” on page 159.

About MSDP performance

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

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

See [“About MSDP server requirements”](#) on page 28.

Table 2-9 MSDP job load performance for an MSDP storage server

When	Description
Normal operation	<p>Normal operation is when all clients have been backed up once.</p> <p>Approximately 15 to 20 jobs can run concurrently and with high performance under the following conditions:</p> <ul style="list-style-type: none">■ 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.

Table 2-9 MSDP job load performance for an MSDP storage server
(continued)

When	Description
Storage approaches full capacity	<p>NetBackup maintains the same number of concurrent backup jobs as during normal operation under the following conditions:</p> <ul style="list-style-type: none">■ The hardware meets minimum requirements. (More capable hardware improves performance.)■ The amount of data that is stored is between 85% to 90% of the capacity of the storage. <p>However, the average time to complete the jobs increases significantly.</p>

How file size may affect the 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.

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

Oracle stream handler

The Oracle stream handler is enabled by default for newly created Oracle policies after an upgrade to NetBackup 8.3. By default, the Oracle stream handler only supports stream-based backups. You can enable and disable the Oracle stream handler per policy or all policies at once using the `cacontrol` command line utility.

When the Oracle stream handler is used, the `FILESERSET` variable is more flexible. You can set `FILESERSET > 1` in the policy or the `bp.conf` without a large decrease in deduplication rates.

The `Marker Entry` file is used to override the default behavior by using the `cacontrol` command utility with the `--sth` flag at a client and or policy level only.

Note: When you use the Oracle stream handler, it is not recommended to use variable-length deduplication.

Update a `Marker Entry` policy by using the following `cacontrol` options:

- You can create and or update a `Marker Entry` policy.

```
cacontrol --sth update  
<STHTYPE> <CLIENT> <POLICY> [SPAUSER] <enabled | disabled>
```

- You can query the setting for the stream handler per policy.

```
cacontrol --sth get <STHTYPE> <CLIENT> <POLICY> [SPAUSER]
```

- You can delete the settings to use the default behavior.

```
cacontrol --sth delete <STHTYPE> <CLIENT> <POLICY> [SPAUSER]
```

Microsoft SQL Server stream handler

You can apply the Microsoft SQL Server stream handler to all of the Microsoft SQL Server version and Azure SQL Server. You can use the **MS-SQL** policy or the **Standard** policy to enable this feature.

You can enable and disable the Microsoft SQL Server stream handler per policy or all policies at once using the `cacontrol` command line utility.

The `Marker Entry` file is used to override the default behavior by using the `cacontrol` command utility with the `--sth` flag at a client and or policy level only.

The `Marker Entry` configuration file is stored at the following location:

```
/MDSP_SERVER/databases/spa/marker.cfg
```

Update a `Marker Entry` policy by using the following `cacontrol` options:

- You can create and or update a `Marker Entry` policy.

```
cacontrol --sth update <STHTYPE> <CLIENT> <POLICY> [SPAUSER]  
<enabled | disabled>
```

- You can query the setting for the stream handler per policy.

```
cacontrol --sth get <STHTYPE> <CLIENT> <POLICY> [SPAUSER]
```

- You can delete the settings to use the default behavior.

```
cacontrol --sth delete <STHTYPE> <CLIENT> <POLICY> [SPAUSER]
```

When you enable the Microsoft SQL Server stream handler, the **Job Details** tab in the Administration Console displays the following:

MS-SQL stream handler enabled

MSDP deployment best practices

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

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

Use fully qualified domain names

Veritas 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 425.

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.

Veritas 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.

Veritas 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

Veritas 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 38.

Introduce MSDP load balancing servers gradually

Veritas 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 38.

Because of the various factors, Veritas 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 127.

See “[About MSDP encryption](#)” on page 129.

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.veritas.com/docs/TECH77575>

About storage unit groups for MSDP

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

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

The best storage savings occur when a backup policy stores its data in the same deduplication destination disk pool instead of across multiple disk pools. For this

reason, the **Failover** method for the **Storage unit selection** uses the least amount of storage. All of the other methods are designed to use different storage every time the backup runs. Veritas recommends that you select the **Failover** method for the **Storage unit selection** type.

Table 2-10 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 storage units with other storage types.
Limitations	NetBackup does not support the following for storage unit groups: <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 141.■ 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 36.

About protecting the MSDP data

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

- Use NetBackup optimized duplication to copy the images to another deduplication node at an off-site location.
Optimized duplication copies the primary backup data to another deduplication pool. It provides the easiest, most efficient method to copy data off-site yet remain in the same NetBackup domain. You then can recover from a disaster that destroys the storage on which the primary copies reside by retrieving images from the other deduplication pool.
See [“Configuring MSDP optimized duplication within the same NetBackup domain”](#) on page 149.
- Use NetBackup replication to copy the deduplicated data to another NetBackup domain off-site.
See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 158.

Veritas also recommends that you back up the MSDP catalog.

See [“About protecting the MSDP catalog”](#) on page 233.

Save the MSDP storage server configuration

Veritas 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 227.

See [“Saving the MSDP storage server configuration”](#) on page 228.

See [“Editing an MSDP storage server configuration file”](#) on page 229.

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, Veritas recommends that you disable the hardware caches. Read and write performance may decline, but you help to avoid data loss.

Provisioning the storage

This chapter includes the following topics:

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

About provisioning the storage for MSDP

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

Provision the storage as follows:

Up to 64 TBs	See Up to 64 TBs of storage .
250 TBs	See 250-TBs of storage .

How many storage instances you provision depends on your storage requirements for your backups. If your requirements are greater than one deduplication node can accommodate, you can configure more than one node.

See [“About MSDP deduplication nodes”](#) on page 19.

Optimized duplication and replication also can affect the number of nodes you provision.

See [“About MSDP optimized duplication and replication”](#) on page 37.

Other NetBackup requirements may affect how you provision the storage.

See [“About MSDP storage and connectivity requirements”](#) on page 21.

How to provision the storage is beyond the scope of the NetBackup documentation. Consult the storage vendor's documentation.

See [“About the NetBackup deduplication destinations”](#) on page 19.

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

Up to 64 TBs of storage

Provision the backup storage so that it appears as a single mount point to the operating system.

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

If you use a separate disk volume for the deduplication database, provision a 1-TB volume on a different mount point than the backup data storage.

250-TBs of storage

NetBackup supports 250 TBs of storage in a single **Media Server Deduplication Pool** on certain operating systems.

See [“About MSDP storage capacity”](#) on page 20.

Before you configure the MSDP storage server, you must provision the volumes. Each volume must conform to the following items:

- Formatted with a file system that NetBackup supports for MSDP. The same file system must be used for all volumes.
- Reside on a separate disk from the other volumes that you allocate for the MSDP storage.
- Mounted on a separate mount point on the computer that you want to use as the MSDP storage server
Veritas recommends that you use a descriptive naming convention for the mount point names.

Steps to configure the 250 TB MSDP using 32 TB volumes

- 1 Create, format, and mount 9 new file systems - one file system must have 1 TB storage space and the other 8 file systems must have 32 TB storage space each.
- 2 Mount the 1 TB file system at `/msdp/cat` and the 32 TB file systems on `/msdp/vol0`, `/msdp/vol1` and so on until each volume is mounted.
- 3 Create a touch a file `/etc/nbapp-release` if it does not exist.

- 4 Create a sub-directory named **data** under each mounted volume. For example, `/msdp/vol0/data`, `/msdp/vol1/data`, `/msdp/vol2/data`, and so on.
- 5 Configure MSDP through the **Storage Server Configuration Wizard**. Ensure that the **Use alternate path for deduplication database** option is selected. Provide the storage path as `/msdp/vol0/data` and the database path as `/msdp/cat`.

- 6 Add additional 32 TB file systems to the deduplication pool:

```
/usr/opensv/pdde/pdcr/bin/crcontrol --dsaddpartition /msdp/vol1/data  
/usr/opensv/pdde/pdcr/bin/crcontrol --dsaddpartition /msdp/vol2/data  
till volume 07...  
/usr/opensv/pdde/pdcr/bin/crcontrol --dsaddpartition /msdp/vol7/data
```

- 7 Review the following command output to verify the created volumes:

```
/usr/opensv/pdde/pdcr/bin/crcontrol --dsstat 2 | grep Mount  
Mount point count: 7
```

For additional information, refer to the following article:

[How to configure a 250TB Media Server Deduplication Pool \(MSDP\) on Linux and Windows](#)

Do not modify MSDP storage directories and files

Unless you are directed to do so by the NetBackup documentation or by a Veritas 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.

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

See “[Resizing the MSDP storage partition](#)” on page 372.

About volume management for NetBackup MSDP

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

For supported systems, see the InfoScale hardware compatibility list at the Veritas website:

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

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

Licensing deduplication

This chapter includes the following topics:

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

About the MSDP license

NetBackup deduplication is licensed separately from base NetBackup.

The license 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.

If you remove the license 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.

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

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

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

Before you try to install or upgrade to a NetBackup version that supports deduplication, you should determine on which operating systems Veritas supports deduplication. See the NetBackup operating system compatibility list:

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

See “[Licensing NetBackup MSDP](#)” on page 52.

Licensing NetBackup MSDP

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

Enter the license 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** on the NetBackup master server, 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 In the **NetBackup License Key** dialog box, click **Close**.
- 5 Restart all the NetBackup services and daemons.

Configuring deduplication

This chapter includes the following topics:

- [Configuring MSDP server-side deduplication](#)
- [Configuring MSDP client-side deduplication](#)
- [About the MSDP Deduplication Multi-Threaded Agent](#)
- [Configuring the Deduplication Multi-Threaded Agent behavior](#)
- [Configuring deduplication plug-in interaction with the Multi-Threaded Agent](#)
- [About MSDP fingerprinting](#)
- [About the MSDP fingerprint cache](#)
- [Configuring the MSDP fingerprint cache behavior](#)
- [About seeding the MSDP fingerprint cache for remote client deduplication](#)
- [Configuring MSDP fingerprint cache seeding on the client](#)
- [Configuring MSDP fingerprint cache seeding on the storage server](#)
- [Enabling 250-TB support for MSDP](#)
- [About MSDP Encryption using NetBackup KMS service](#)
- [About MSDP Encryption using external KMS server](#)
- [Configuring a storage server for a Media Server Deduplication Pool](#)
- [Configuring a storage server for a PureDisk Deduplication Pool](#)
- [About disk pools for NetBackup deduplication](#)
- [Configuring a disk pool for deduplication](#)

- Creating the data directories for 250-TB MSDP support
- Adding volumes to a 250-TB Media Server Deduplication Pool
- Configuring a Media Server Deduplication Pool storage unit
- Configuring client attributes for MSDP client-side deduplication
- Disabling MSDP client-side deduplication for a client
- About MSDP compression
- About MSDP encryption
- MSDP compression and encryption settings matrix
- Configuring encryption for MSDP backups
- Configuring encryption for MSDP optimized duplication and replication
- About the rolling data conversion mechanism for MSDP
- Modes of rolling data conversion
- MSDP encryption behavior and compatibilities
- Configuring optimized synthetic backups for MSDP
- About a separate network path for MSDP duplication and replication
- Configuring a separate network path for MSDP duplication and replication
- About MSDP optimized duplication within the same domain
- Configuring MSDP optimized duplication within the same NetBackup domain
- About MSDP replication to a different domain
- Configuring MSDP replication to a different NetBackup domain
- About configuring MSDP optimized duplication and replication bandwidth
- About storage lifecycle policies
- About the storage lifecycle policies required for Auto Image Replication
- Creating a storage lifecycle policy
- About MSDP backup policy configuration
- Creating a backup policy
- Resilient Network properties

- Specifying resilient connections
- Adding an MSDP load balancing server
- About variable-length deduplication on NetBackup clients
- About the MSDP pd.conf configuration file
- Editing the MSDP pd.conf file
- About the MSDP contentrouter.cfg file
- About saving the MSDP storage server configuration
- Saving the MSDP storage server configuration
- Editing an MSDP storage server configuration file
- Setting the MSDP storage server configuration
- About the MSDP host configuration file
- Deleting an MSDP host configuration file
- Resetting the MSDP registry
- About protecting the MSDP catalog
- Changing the MSDP shadow catalog path
- Changing the MSDP shadow catalog schedule
- Changing the number of MSDP catalog shadow copies
- Configuring an MSDP catalog backup
- Updating an MSDP catalog backup policy
- About MSDP FIPS compliance
- Configuring the NetBackup client-side deduplication to support multiple interfaces of MSDP
- About MSDP multi-domain support
- About MSDP mutli-domain VLAN Support

Configuring MSDP server-side deduplication

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

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

The *NetBackup Administrator's Guide* describes how to configure a base NetBackup environment.

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

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

Table 5-1 MSDP configuration tasks

Step	Task	Procedure
Step 1	Install the license for deduplication	See "Licensing NetBackup MSDP" on page 52.
Step 2	Create NetBackup log file directories on the master server and the media servers	See "NetBackup MSDP log files" on page 407. See "Creating NetBackup log file directories for MSDP" on page 406.
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 59. 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 the MSDP fingerprint cache" on page 68. See "Configuring the MSDP fingerprint cache behavior" on page 69.
Step 5	Enable support for 250-TB MSDP	Before you configure a storage server that hosts a 250-TB Media Server Deduplication Pool , you must enable support for that size storage. See "Enabling 250-TB support for MSDP" on page 76.

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

Step	Task	Procedure
Step 6	Configure a deduplication storage server	<p>How many storage servers you configure depends on: your storage requirements and on whether or not you use optimized duplication or replication. When you configure a storage server, the wizard also lets you configure a disk pool and a storage unit.</p> <p>See “About MSDP storage servers” on page 27.</p> <p>See “MSDP storage path properties” on page 94.</p> <p>See “About MSDP optimized duplication and replication” on page 37.</p> <p>Which type of storage server to configure depends on the storage destination.</p> <p>See “About the NetBackup deduplication destinations” on page 19.</p> <p>See “Configuring a storage server for a Media Server Deduplication Pool” on page 80.</p>
Step 7	Configure a disk pool	<p>If you already configured a disk pool when you configured the storage server, you can skip this step.</p> <p>How many disk pools you configure depends on: your storage requirements and on whether or not you use optimized duplication or replication.</p> <p>See “About disk pools for NetBackup deduplication” on page 107.</p> <p>See “Configuring a disk pool for deduplication” on page 109.</p>
Step 8	Create the data directories for 250-TB support	<p>For a 250-TB Media Server Deduplication Pool, you must create the data directories under the mount points for the storage directories.</p> <p>See “Creating the data directories for 250-TB MSDP support” on page 118.</p>
Step 9	Add the other volumes for 250-TB support	<p>For a 250-TB Media Server Deduplication Pool, you must add the second and third volumes to the disk pool.</p> <p>See “Adding volumes to a 250-TB Media Server Deduplication Pool” on page 119.</p>
Step 10	Configure a storage unit	See “Configuring a Media Server Deduplication Pool storage unit” on page 121.
Step 11	Enable encryption	<p>Encryption is optional.</p> <p>See “Configuring encryption for MSDP backups” on page 132.</p>
Step 12	Configure optimized synthetic backups	<p>Optimized synthetic backups are optional.</p> <p>See “Configuring optimized synthetic backups for MSDP” on page 138.</p>

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

Step	Task	Procedure
Step 13	Configure MSDP restore behavior	<p>Optionally, you can configure NetBackup to bypass media servers during restores.</p> <p>See “How MSDP restores work” on page 373.</p> <p>See “Configuring MSDP restores directly to a client” on page 374.</p>
Step 14	Configure optimized duplication copy	<p>Optimized duplication is optional.</p> <p>See “About MSDP optimized duplication within the same domain” on page 141.</p>
Step 15	Configure replication	<p>Replication is optional.</p> <p>See “About MSDP replication to a different domain” on page 156.</p>
Step 16	Configure a backup policy	<p>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.</p> <p>See “About MSDP backup policy configuration” on page 199.</p> <p>See “Creating a backup policy” on page 200.</p>
Step 17	Specify advanced deduplication settings	<p>Advanced settings are optional.</p> <p>See “About the MSDP pd.conf configuration file” on page 209.</p> <p>See “Editing the MSDP pd.conf file” on page 210.</p> <p>See “MSDP pd.conf file parameters” on page 210.</p>
Step 18	Protect the MSDP data and catalog	<p>See “About protecting the MSDP data” on page 45.</p> <p>See “About protecting the MSDP catalog” on page 233.</p>

Configuring MSDP client-side deduplication

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

See [“Configuring MSDP server-side deduplication”](#) on page 55.

Table 5-2 Client deduplication configuration tasks

Step	Task	Procedure
Step 1	Configure media server deduplication	See “Configuring MSDP server-side deduplication” on page 55.

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

Step	Task	Procedure
Step 2	Learn about client deduplication	See “About NetBackup Client Direct deduplication” on page 30.
Step 3	Configure a resilient connection for remote office clients	Resilient connections are optional. See “About MSDP remote office client deduplication” on page 32. See “Resilient Network properties” on page 200. See “Specifying resilient connections” on page 204.
Step 4	Enable client-side deduplication	See “Configuring client attributes for MSDP client-side deduplication” on page 125.
Step 5	Configure remote client fingerprint cache seeding	Configuring remote client fingerprint cache seeding is optional. See “Configuring MSDP fingerprint cache seeding on the client” on page 72. See “About seeding the MSDP fingerprint cache for remote client deduplication” on page 70. See “Configuring MSDP fingerprint cache seeding on the storage server” on page 73.
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 restores directly to a client” on page 374.

About the MSDP Deduplication Multi-Threaded Agent

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 any host that deduplicates data: the storage server, load balancing servers, or clients that deduplicate their own data. For each host on which you want to use the Multi-Threaded Agent, you must configure the deduplication plug-in to use it.

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, AIX, 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 ■ <code>SEGKSIZE</code> is greater than 128 (<code>pd.conf</code> file) ■ <code>DONT_SEGMENT_TYPES</code> enabled (<code>pd.conf</code> file) ■ <code>MATCH_PDRO = 1</code> (<code>pd.conf</code> file) <p>See “MSDP <code>pd.conf</code> file parameters” on page 210.</p>

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

Item	Description
Policy-based compression or encryption	<p>If NetBackup policy-based compression or encryption is enabled on the backup policy, NetBackup does not use the Deduplication Multi-Threaded Agent.</p> <p>Veritas 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 127.</p> <p>See “About MSDP encryption” on page 129.</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 59.

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)

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

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

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

Process Parameter	Description
MaxConcurrentSessions	<p>The maximum number of concurrent sessions that the Multi-Threaded Agent processes. If it receives a backup job when the <code>MaxConcurrentSessions</code> value is reached, the job runs as a single-threaded job.</p> <p>By default, the deduplication plug-in sends backup jobs to the Multi-Threaded Agent on a first-in, first-out basis. However, you can configure which clients and which backup policies the deduplication plug-in sends to the Multi-Threaded Agent. The <code>MTSTRM_BACKUP_CLIENTS</code> and <code>MTSTRM_BACKUP_POLICIES</code> parameters in the <code>pd.conf</code> control the behavior. Filtering the backup jobs that are sent to the Multi-Threaded Agent can be very helpful on the systems that have many concurrent backup jobs.</p> <p>See “MSDP pd.conf file parameters” on page 210.</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 pd.conf file parameters” on page 210.</p> <p>Possible values: 1-32, inclusive.</p> <p>Warning: Veritas 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*)

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

Threads parameters

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

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

Threads Parameter	Description
<code>BackupFpThreads</code>	<p>The number of threads to use per session to fingerprint incoming data.</p> <p>Default value: <code>BackupFpThreads</code>= (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 210.</p>
<code>BackupSendThreads</code>	<p>The number of threads to use per session to send data to the storage server during a backup operation.</p> <p>Default value: <code>BackupSendThreads</code>=1 for servers and <code>BackupSendThreads</code>=2 for clients.</p> <p>Possible values: 1-32, inclusive.</p>
<code>MaintenanceThreadPeriod</code>	<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 209.

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 described in another topic.
- See [“MSDP pd.conf file parameters”](#) on page 210.
- 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.

NetBackup 8.0 and previous versions use fingerprinting calculations that are based on the MD5-like algorithm. Starting with NetBackup 8.1, the fingerprinting calculations are based on a more secure SHA-2 algorithm. On a system that is upgraded to the 8.1 version, every new segment is computed with the SHA-2 algorithm. A data rolling conversion task works in the background to convert the existing MD5-like fingerprints to SHA-2 fingerprints, gradually.

See [“About the rolling data conversion mechanism for MSDP”](#) on page 134.

NetBackup 8.1 can handle both fingerprint types, and the new server is compatible with old clients and old servers. When you perform a backup from an old client to a new server or when you duplicate data from an old server to a new server, conversion from MD5-like to SHA-2 occurs inline on the new server before the data is saved to the disk. Similarly, when you duplicate data from a new server to an old server, conversion from SHA-2 to MD5-like occurs inline on the new server before the data is sent to the old server.

Notes and restrictions that there are some known issues for the compatibility support.

- The fingerprint conversion requires additional computation time. The interaction between old clients and old servers and new server is slower than if both the client and the server are new.
- You cannot restore data that is backed up using SHA-2 algorithm on a media server that uses the MD5-like algorithm. However, you may choose to restore the SHA-2 fingerprint data on a new media server.
- Similarly, you cannot use client-direct restore to restore data that is backed up using Client Direct deduplication on a media server that uses the MD5-like algorithm. However, you may choose to restore the data on a new media server.
- If you are using two types of media servers for load balancing, where one media server uses MD5-like algorithm and the other media server uses the SHA-2 algorithm, the initial backup may lose deduplication ratio. Therefore, split the old media servers and the new media servers into different groups, and create different storage unit for each of them.
- When data is backed up from a NetBackup 7.5 or previous version client, using Client Direct deduplication, most of the data is transferred over the network and deduplicated on the server. This may save storage, but it does not reduce network throughput. It is recommended that you upgrade the NetBackup client to the latest version.

See [“About the MSDP fingerprint cache”](#) on page 68.

See [“Media server deduplication backup process”](#) on page 393.

See [“MSDP client-side deduplication backup process”](#) on page 396.

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 69.

You can also control the fingerprint cache seeding for clients.

See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 70.

Configuring the MSDP fingerprint cache behavior

You can configure the cache loading behavior.

See [“About the MSDP fingerprint cache”](#) on page 68.

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 69.

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 227.

Table 5-8 Cache load parameters

Behavior	Description
<code>CacheLoadThreadNum</code>	<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>
<code>MaxCacheSize</code>	<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 50%.</p>

About seeding the MSDP fingerprint cache for remote client deduplication

Veritas 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 “Configuring MSDP fingerprint cache seeding on the client” on page 72.
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 “Configuring MSDP fingerprint cache seeding on the storage server” on page 73.

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

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

If these two conditions are not met, NetBackup may use multiple streams. If the **Attributes** tab **Limit jobs per policy** is set to a number less than the total number

of streams, only those streams use the seeded images to populate the cache. Any streams that are greater than the **Limit jobs per policy** value do not benefit from seeding, and their cache hit rates may be close to 0%.

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

The following items are example of informational messages that show that seeding occurred:

Activity Monitor Job Details

```
1/2/2015 2:18:23 AM - Info nbmaster1(pid=6340)
StorageServer=PureDisk:nbmaster1; Report=PDDO
Stats for (nbmaster1): scanned: 3762443 KB, CR
sent: 1022 KB, CR sent over FC: 0 KB, dedup:
100.0%, cache hits: 34364 (100.0%)

1/2/2015 2:18:24 AM - Info nbmaster1(pid=6340)
StorageServer=PureDisk:nbmaster1; Report=PDDO
Stats for (nbmaster1): scanned: 1 KB, CR sent:
0 KB, CR sent over FC: 0 KB, dedup: 100.0%
```

Deduplication plug-in log (pdplugin.log) on the client

```
01/02/15 02:15:17 [4452] [4884] [DEBUG] PDSTS:
cache_util_get_cache_dir: enter
db=/nbmaster1#1/2, scp='', bc=opscenter1,
bp=seedfinal, bl=4096

01/02/15 02:15:17 [4452] [4884] [DEBUG] PDSTS:
cache_util_get_cache_dir: new backup, using
existing client seeding directory

01/02/15 02:15:17 [4452] [4884] [DEBUG] PDSTS:
cache_util_get_cache_dir: exit
db=/nbmaster1#1/2, scp='', bc=opscenter1,
bp=seedfinal, bl=4096,
cachedir_buf='/nbmaster1#1/2/#pdseed/opscenter1'
err=0
```

See ["NetBackup MSDP log files"](#) on page 407.

```

Deduplication proxy server log (nbostpxy.log) on the client 02:15:17.417[4452.4884] [DEBUG] [dummy] [11:bptm:6340:nbraster1] [DEBUG]
PDSTS: cache_util_get_cache_dir: enter
db=/nbraster1#1/2, scp='', bc=opscenter1,
bp=seedfinal, bl=4096

02:15:17.433[4452.4884] [DEBUG] [dummy] [11:bptm:6340:nbraster1] [DEBUG]
PDSTS: cache_util_load_fp_cache_nbu: enter
dir_path=/nbraster1#1/2/#pdseed/opscenter1,
t=16s, me=1024

02:15:17.449[4452.4884] [DEBUG] [dummy] [11:bptm:6340:nbraster1] [DEBUG]
PDSTS: cache_util_load_fp_cache_nbu: adding
'nbraster1_1420181254_C1_F1.img' to cache list
(1)

02:15:17.449[4452.4884] [DEBUG] [dummy] [11:bptm:6340:nbraster1] [DEBUG]
PDSTS: cache_util_load_fp_cache_nbu: opening
/nbraster1#1/2/#pdseed/opscenter1/nbraster1_1420181254_C1_F1.img
for image cache (1/1)

02:15:29.585[4452.4884] [DEBUG] [dummy] [11:bptm:6340:nbraster1] [DEBUG]
PDVFS: pdvfs_lib_log: soRead: segment
c32b0756d491871c45c71f811fbd73af already
present in cache.

02:15:29.601[4452.4884] [DEBUG] [dummy] [11:bptm:6340:nbraster1] [DEBUG]
PDVFS: pdvfs_lib_log: soRead: segment
346596a699bd5f0ba5389d4335bc7429 already
present in cache.

```

See “[NetBackup MSDP log files](#)” on page 407.

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

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

See “[About the MSDP fingerprint cache](#)” on page 68.

Configuring MSDP fingerprint cache seeding on the client

Seeding on the client requires the following:

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

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

See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 70.

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

To seed the MSDP fingerprint cache on the client

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

Specify the setting in the following format:

clienthostmachine,backuppolicy,date

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

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

backuppolicy The backup policy for that client.

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

See [“Editing the MSDP pd.conf file”](#) on page 210.

See [“MSDP pd.conf file parameters”](#) on page 210.

Configuring MSDP fingerprint cache seeding on the storage server

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

database_path/databases/catalog/2/#pdseed/client_name

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

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

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

See [“About seeding the MSDP fingerprint cache for remote client deduplication”](#) on page 70.

To seed the fingerprint cache from the storage server

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

UNIX: `/usr/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`

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

See [“NetBackup seedutil options”](#) on page 75.

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

```
seedutil -list_clients
```

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

```
seedutil -clear client_name
```

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

NetBackup seedutil options

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

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

The following items describe the options:

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

The following are the directories in which the command resides:

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

Enabling 250-TB support for MSDP

Before you configure a storage server for a 250-TB **Media Server Deduplication Pool**, you must enable support for the multiple volumes that are required.

See [“About MSDP storage capacity”](#) on page 20.

See [“About provisioning the storage for MSDP”](#) on page 47.

To enable support for a 250-TB NetBackup MSDP

- ◆ On the NetBackup media server that you want to use as the MSDP storage server, create an empty `nbapp-release` file in the `/etc` directory. The following is an example that shows how to use the `touch` command to create the file:

```
touch /etc/nbapp-release
```

Use the following command to create the file on Windows:

```
mkdir c:\etc  
echo Windows_BYO > "c:\etc\nbapp-release"
```

For additional information, refer to the following article:

[How to configure a 250TB Media Server Deduplication Pool \(MSDP\) on Linux and Windows](#)

About MSDP Encryption using NetBackup KMS service

NetBackup incorporates Key Management Server (KMS) with Media Server Deduplication Pool.

MSDP encryption carries out segment-level encryption and assigns a unique encryption key for every data segment. A customer key is retrieved from NetBackup KMS to encrypt the segment key.

You can configure the KMS service from the NetBackup Administration Console or the NetBackup command line during storage server configuration.

Note: You cannot disable the MSDP KMS service once you enable it.

If the KMS service is not available for MSDP or the key in the KMS service that is used by MSDP is not available, then MSDP waits in an infinite loop. When MSDP goes in an infinite loop, few commands that you run might not respond.

After you configure KMS encryption or once the MSDP processes restart, check the KMS encryption status after the first backup finishes.

The keys in the key dictionary must not be deleted, deprecated, or terminated.

You can use the following commands to get the status of the KMS mode:

- For UNIX:

```
/usr/opensv/pdde/pdcr/bin/crcontrol --getmode
```

For MSDP cloud, run the following `keydictutil` command to check if the LSU is in KMS mode:

```
/usr/opensv/pdde/pdcr/bin/keydictutil --list
```

- For Windows:

```
<install_path>\Veritas\pdde\crcontrol.exe --getmode
```

Note: If you use the `nbdevconfig` command to add a new encrypted cloud Logical Storage Unit (LSU) and an encrypted LSU exists in this MSDP, the `keygroupname` must be the same as the `keygroupname` in the previous encrypted LSU.

For enabling KMS, refer to the following topics:

See [“Configuring a storage server for a Media Server Deduplication Pool”](#) on page 80.

Upgrading KMS for MSDP

Before you upgrade KMS encryption from NetBackup version earlier than 8.1.1, complete the following steps. During the NetBackup upgrade, KMS rolling conversion runs along with MSDP encryption rolling conversion.

For NetBackup version earlier than 8.1.1, the supported NetBackup upgrade paths are:

- NetBackup 7.7.3 to 8.1.2 or later
- NetBackup 8.0 to 8.1.1 or later
- NetBackup 8.1 to 8.1.1 or later

For additional information, refer to the *Configuring KMS* section in the *Veritas NetBackup Security and Encryption Guide*.

Before you upgrade KMS, complete the following steps:

Note: The following steps are not supported on Solaris OS. For Solaris, refer to the following article:

[Upgrade KMS encryption for MSDP on the Solaris platform](#)

1 Create an empty database using the following command:

- For UNIX:

```
/usr/opensv/netbackup/bin/nbkms -createemptydb
```

- For Windows:

```
<install_path>\Veritas\NetBackup\bin\nbkms.exe -createemptydb
```

Enter the following parameters when you receive a prompt:

- Enter the HMK passphrase
Enter a password that you want to set as the host master key (HMK) passphrase. Press Enter to use a randomly generated HMK passphrase. The passphrase is not displayed on the screen.
- Enter HMK ID
Enter a unique ID to associate with the host master key. This ID helps you to determine an HMK associated with any key store.
- Enter KPK passphrase
Enter a password that you want to set as the key protection key (KPK) passphrase. Press Enter to use a randomly generated HMK passphrase. The passphrase is not displayed on the screen.
- Enter KPK ID
Enter a unique ID to associate with the key protection key. This ID helps you to determine a KPK associated with any key store.

After the operation completes successfully, run the following command on the master server to start KMS:

- For UNIX:

```
/usr/opensv/netbackup/bin/nbkms
```

- For Windows:

```
sc start NetBackup Key Management Service
```

2 Create a key group and an active key by entering the following commands:

- For UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbkmsutil -createkg -kgname  
msdp
```

```
/usr/opensv/netbackup/bin/admincmd/nbkmsutil -createkey -kname  
msdp -keyname name -activate
```

- For Windows:

```
<install_path>\Veritas\NetBackup\bin\admincmd\nbkmsutil.exe  
-createkg -kname msdp  
<install_path>\Veritas\NetBackup\bin\admincmd\nbkmsutil.exe  
-createkey -kname msdp -keyname name -activate
```

Enter a password that you want set as the key passphrase.

- 3 Create a `kms.cfg` configuration file at the following location on the NetBackup media server where you have configured the MSDP storage:

- On UNIX:

```
/usr/opensv/pdde/kms.cfg
```

- On Windows:

```
<install_path>\Veritas\pdde\kms.cfg
```

Add the following content to the `kms.cfg` file:

```
[KMSOptions]  
KMSEnable=true  
KMSKeyGroupName=YourKMSKeyGroupName  
KMSServerName=YourKMSServerName  
KMSType=0
```

For `KMSServerName`, enter the hostname of the server where the KMS service runs, mainly the master server hostname.

After completing the steps, you can upgrade MSDP.

About MSDP Encryption using external KMS server

NetBackup supports keys from an external key management service (KMS) server that encrypts data in case of MSDP storage. Keys are retrieved from the external KMS server to encrypt the backup data.

For information about external KMS support, see the [NetBackup Security and Encryption Guide](#).

The other information remains the same as mentioned in the following topic:

See [“About MSDP Encryption using NetBackup KMS service”](#) on page 76.

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 27.

The type of storage.	Select Media Server Deduplication Pool for the type of disk storage.
The credentials for the deduplication engine.	See “About the NetBackup Deduplication Engine credentials” on page 34.
The storage paths.	See “MSDP storage path properties” on page 94.
The network interface.	See “About the network interface for MSDP” on page 35.
The load balancing servers, if any.	See “About MSDP storage servers” on page 27.

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

Prerequisite	For a 96-TB Media Server Deduplication Pool , you must create the required directories before you configure the storage server. See “Creating the data directories for 250-TB MSDP support” on page 118.
--------------	--

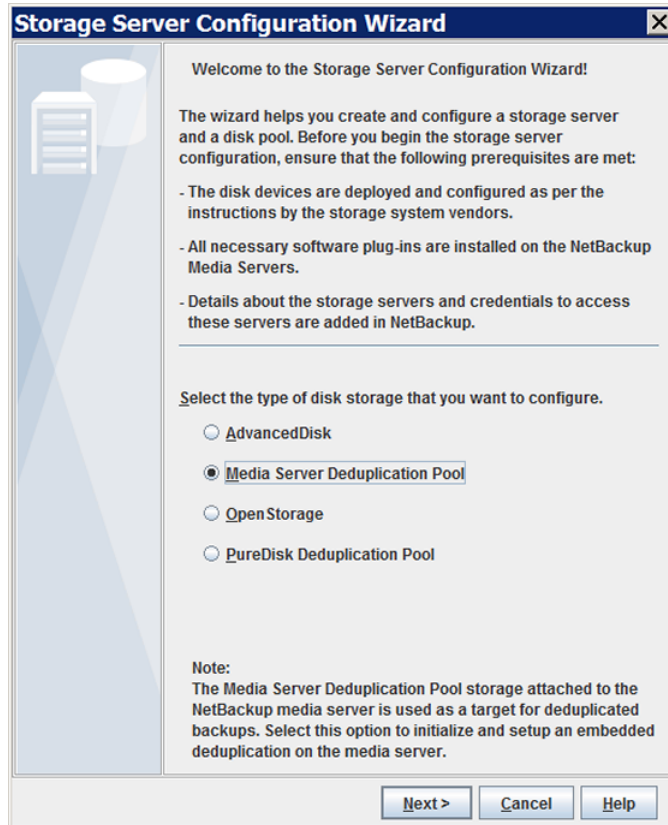
To configure a NetBackup storage server for a Media Server Deduplication Pool

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

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

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

The following is an example of the wizard panel:



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

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

Storage Server Configuration Wizard

Add Storage Server

Provide storage server details.

Select the media server that connects to the storage. The media server runs the core NetBackup Deduplication Engine components and functions as the storage server.

Media server:

Storage server type:

Media Server Deduplication Pool

Storage server name:

Define credentials

User name:

Password:

Confirm password:

< Back

Next >

Cancel

Help

Media server	<p>Select the media server that you want to configure as the storage server.</p> <p>You can add deduplication load balancing servers on the next wizard panel.</p>
Username	<p>Enter the user name for the NetBackup Deduplication Engine.</p> <p>See “About the NetBackup Deduplication Engine credentials” on page 34.</p>
Password	<p>Enter the password for the NetBackup Deduplication Engine.</p>
Confirm password	<p>To confirm the password, re-enter the password.</p>

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

The **Storage Server Properties** panel appears.

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

See “MSDP storage path properties” on page 94.

See “MSDP network interface properties” on page 97.


The following is an example of the wizard panel:

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

No media servers installed	The Storage Server Encryption panel appears. Go to step 8.
Media servers are installed.	The Deduplication Load Balancing panel appears. Continue to step 6.

- 6 On the **Deduplication Load Balancing** panel, select other NetBackup media servers to use for deduplication. Selecting load balancing servers is optional. See [“About MSDP load balancing servers”](#) on page 28.

The following is an example of the wizard panel:



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

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

Note: Load balancing for Cloud Catalyst storage servers or Cloud Catalyst appliances is supported. See the following topic for configuration considerations. See [“Configuring source control or target control optimized duplication for Cloud Catalyst”](#) on page 277.

- 7 On the **Storage Server Encryption** panel, you can enable encryption and KMS for the Media Server Deduplication Pool. The **Enable KMS** option is available when you select **Enable Encryption**.

If none of the KMS servers is configured, NetBackup KMS is configured:

The following is an example of the wizard panel:

Storage Server Configuration Wizard

Storage Server Encryption

Provide storage server encryption.

☒ Enable Encryption

☒ Enable KMS

Key Management Server (KMS) is not configured. This will configure KMS server to enable data encryption.

KMS server name:ray-win-debug

Host master key (HMK) passphrase:

Re-enter HMK passphrase:

Host master key id:

Key protection key (KPK) passphrase:


Re-enter KPK passphrase:

Key protection key id:

Key name:

Key passphrase:

Re-enter Key passphrase:

 Once you define the storage server details on this screen, you cannot modify them. For more information, click Help.

< Back

Next >

Cancel

Help

- See [“About MSDP encryption”](#) on page 129.
- See [“About MSDP Encryption using NetBackup KMS service”](#) on page 76.
- See [“About MSDP Encryption using external KMS server”](#) on page 79.
- If you select **Enable KMS** for the first time, as a one-time KMS configuration, you must enter the following information:

Option	Description
KMS server name	The name of the Key Management Server.
Host master key (HMK) passphrase	Enter a password that you want to set as the host master key (HMK) passphrase.

Option	Description
Host master key id	Enter a unique ID to associate with the host master key. This ID helps you to determine an HMK associated with any key store.
Key protection key (KPK) passphrase	Enter a password that you want to set as the key protection key (KPK) passphrase. This ID helps you to determine a KPK associated with any key store.
Key protection key id	Enter a unique ID to associate with the key protection key.
Key name	Enter a name of the key.
Key passphrase	Enter a password that you want set as the key passphrase.

- If you select **Enable KMS** and if NetBackup KMS is already configured on the Master Server, you must enter the following information:

Option	Description
Key name	Enter a name of the key.
Key passphrase	Enter a password that you want set as the key passphrase.

- If you select **Enable KMS** and if external KMS is already configured on the Master Server, you must enter the following information:

Option	Description
Key group name	Enter a name of the key group.
<p>Ensure that a key with a custom attribute set with value of key group name already exists in the external KMS server.</p> <p>For more information about KMS configuration, see the NetBackup Security and Encryption Guide.</p>	

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

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

The following is an example of the wizard panel:


The screenshot shows a window titled "Storage Server Configuration Wizard" with a subtitle "Storage Server Configuration Summary" and the instruction "Verify the storage server configuration." Below this, a message states: "Review the storage server configuration summary. You can change the configuration, if required." A table lists the configuration details:

Storage server type:	Media Server Deduplication Pool
Storage server name:	msdp-ss-host.example.com
Media server name:	msdp-ss-host.example.com
User name:	backup_administrator
Storage path:	/storage_msdp
Database path:	/storage_msdp_cat
Network interface:	Use Default
Deduplication load balancing media serv...	msdp-lb1-host.example.com msdp-lb2-host.example.com

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

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

- 9 The **Storage Server Configuration Status** wizard panel describes the status of the operation.



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

- Configure a disk pool

Ensure that **Create a disk pool using the storage server that you have just created** is selected and then click **Next**.

The **Select Volumes** panel appears. Continue to the next step.
- Exit

Click **Close**.

You can configure a disk pool at another time.

If storage server creation fails, troubleshooting help is available. See [“Troubleshooting MSDP configuration issues”](#) on page 414.

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

The following is an example of the wizard panel:

Storage Server Configuration Wizard

Volume Selection
Select volumes to use in the disk pool.

Storage server: msdp-ss-host.example.com

Storage server type: PureDisk

Select storage server volumes to add to the disk pool.

Volume Name	Available Space	Raw Size	Replication
<input checked="" type="checkbox"/> PureDiskVolume	22.21 GB	22.25 GB	None

Disk Pool Size

Total available space: 22.21 GB

Total raw size: 22.25 GB

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


< Back Next > Cancel Help

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

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

See “[Media Server Deduplication Pool properties](#)” on page 116.

The following is an example of the wizard panel:



Storage Server Configuration Wizard

Additional Disk Pool Information
 Provide additional disk pool information.

Storage server: msdp-ss-host.example.com

Storage server type: PureDisk

Disk Pool Size

Total available space: 22.21 GB

Total raw size: 22.25 GB

Disk Pool name:

Comments:

High water mark: %

Low water mark: %

Maximum I/O Streams

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

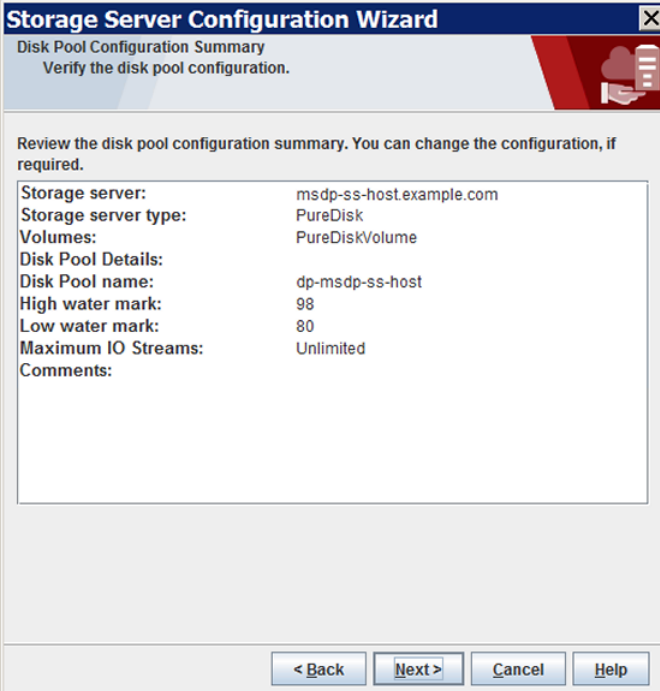
☐ Limit I/O streams: per volume

< Back Next > Cancel Help

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

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

The following is an example of the wizard panel:



Storage Server Configuration Wizard

Disk Pool Configuration Summary
 Verify the disk pool configuration.


Review the disk pool configuration summary. You can change the configuration, if required.

Storage server:	msdp-ss-host.example.com
Storage server type:	PureDisk
Volumes:	PureDiskVolume
Disk Pool Details:	
Disk Pool name:	dp-msdp-ss-host
High water mark:	98
Low water mark:	80
Maximum IO Streams:	Unlimited
Comments:	

< Back Next > Cancel Help

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

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



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

- Configure a storage unit

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

Click **Close**.

You can configure one or more storage units later.

See [“Configuring a Media Server Deduplication Pool storage unit”](#) on page 121.

- 14 Enter the appropriate information for the storage unit.
- See “Media Server Deduplication Pool storage unit properties” on page 123.
- The following is an example of the wizard panel:

Storage Server Configuration Wizard

Storage Unit Creation
Enter details to create storage unit.

Disk pool:

dp-msdp-ss-host

Storage server type:

PureDisk

Storage unit name:

stu-dp-msdp-ss-host

Media Server

☒ Use any available media server to transport data

☐ Only use the selected media servers:

Media Servers

☐ msdp-lb1-host.example.com

☐ msdp-lb2-host.example.com

Maximum concurrent jobs:

1

Maximum fragment size:

51200

Megabytes

< Back

Next >

Cancel

Help

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

MSDP storage path properties

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

Table 5-10 MSDP storage path properties

Property	Description
Storage path	<p>The path to the storage. The storage path is the directory in which NetBackup stores the raw backup data. Backup data should not be stored on the system disk.</p> <p>Because the storage requires a directory path, do not use only the root node (/) or drive letter (E:\) as the storage path. (That is, do not mount the storage as a root node (/) or a drive letter (E:\).</p> <p>For a 250-TB Media Server Deduplication Pool, you must enter the path name of the mount point for the volume that you consider the first 32-TB storage volume. The following is an example of a volume naming convention for the mount points for the backups:</p> <pre>/msdp/vol0 <--- The first volume /msdp/vol1 /msdp/vol2</pre> <p>NetBackup supports 250-TB deduplication pools on a subset of supported systems.</p> <p>See “About MSDP storage capacity” on page 20.</p> <p>See “About provisioning the storage for MSDP” on page 47.</p> <p>See “Creating the data directories for 250-TB MSDP support” on page 118.</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, G:\MSDP_Storage) <p>NetBackup requirements for the deduplication storage paths may affect how you expose the storage.</p> <p>See “About MSDP storage and connectivity requirements” on page 21.</p>

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

Property	Description
Use alternate path for deduplication database	<p>By default, NetBackup uses the storage path for the MSDP database (that is, the MSDP catalog) location. The MSDP database is different than the NetBackup catalog.</p> <p>Select this option to use a location other than the default for the deduplication database.</p> <p>For a 250-TB Media Server Deduplication Pool, you must select this option.</p> <p>See “About provisioning the storage for MSDP” on page 47.</p> <p>For performance optimization, it is recommended that you use a separate disk volume for the deduplication database than for the backup data.</p>
Database path	<p>If you selected Use alternate path for deduplication database, enter the path name for the database. The database should not be stored on the system disk.</p> <p>For a 250-TB Media Server Deduplication Pool, you must enter the path name of the partition that you created for the MSDP catalog. For example, if the naming convention for your mount points is <code>/msdp/volx</code>, the following path is recommended for the MSDP catalog directory:</p> <pre>/msdp/cat</pre> <p>See “About provisioning the storage for MSDP” on page 47.</p> <p>For performance optimization, it is recommended that you use a separate disk volume for the deduplication database than for the backup data.</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: <code>_ - : . / \</code> Windows: <code>_ - : . \</code> (a colon (<code>:</code>) is allowed only after a drive letter (for example, <code>F:\MSDP_Storage</code>)

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

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

MSDP network interface properties

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

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

Table 5-11 MSDP network interface properties

Property	Description
Use specific network interface	Select this option to specify a network interface for the deduplication traffic. If you do not specify a network interface, NetBackup uses the operating system host name value. See “About the network interface for MSDP” on page 35.
Interface	If you selected Use specific network interface , enter the interface name.

Configuring a storage server for a PureDisk Deduplication Pool

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

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

See [“About the NetBackup deduplication destinations”](#) on page 19.

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

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

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

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

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

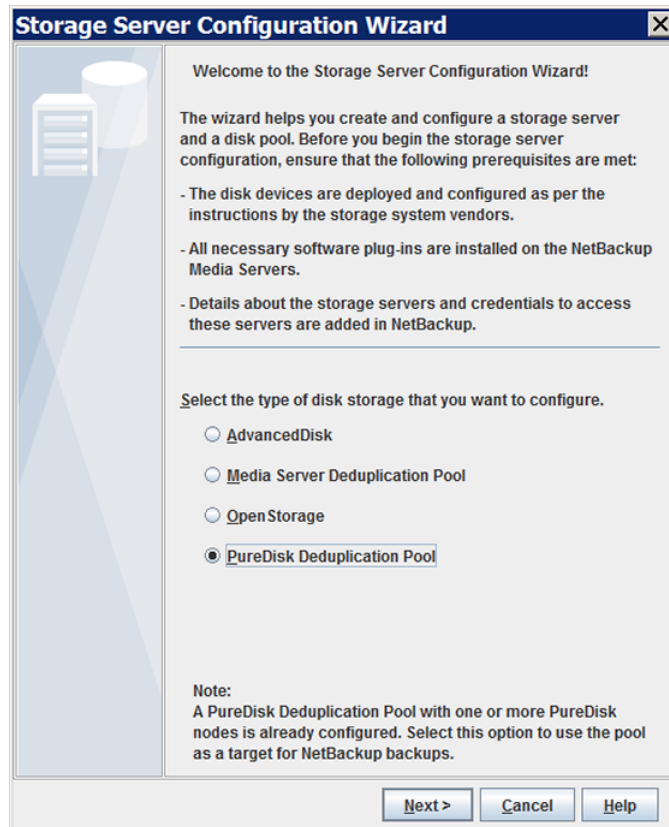
- 2 In the right pane, click **Configure Disk Storage Servers**.

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

- 3 On the **Welcome** panel, select **PureDisk Deduplication Pool**.

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

The following is an example of the wizard panel:



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

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

Storage Server Configuration Wizard

Add Storage Server
Provide storage server details.

Select a media server to use to query the PureDisk Storage Pool Authority host for its capabilities. This media server is also used for deduplication

Media server:

Storage server type:

Storage server name:

Enter SPA credentials

User name:

Password:

Confirm password:

< Back Next > Cancel Help

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

You can add other deduplication load balancing servers on the next wizard panel. |
| Storage server name | Enter the host name of the NetBackup 5000 series appliance in the Storage server name field. |
| Username | Enter the user name for the NetBackup 5000 series appliance. |
| Password | Enter the password for the NetBackup 5000 series appliance. |
| Confirm password | To confirm the password, re-enter the password. |

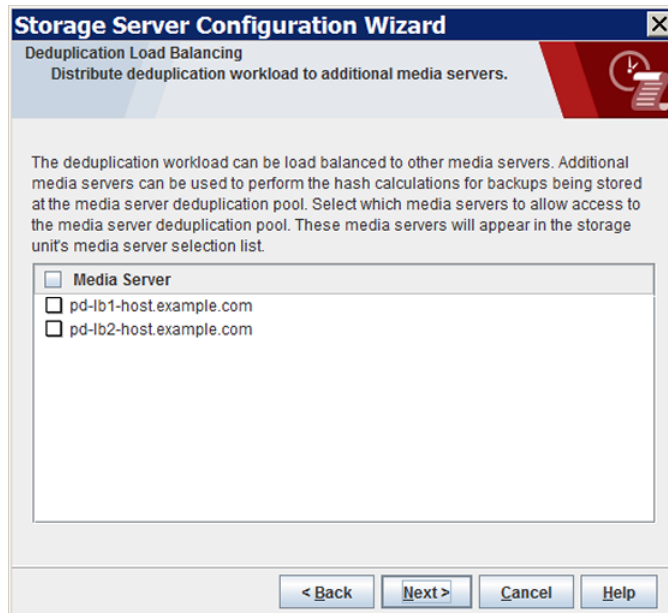
After you enter the information, click **Next**.

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

- 5 On the **Deduplication Load Balancing** panel select other NetBackup media servers to use for deduplication load balancing. The wizard populates the media servers that you select with the Storage Pool Authority credentials you entered on a previous panel. Selecting load balancing servers is optional.

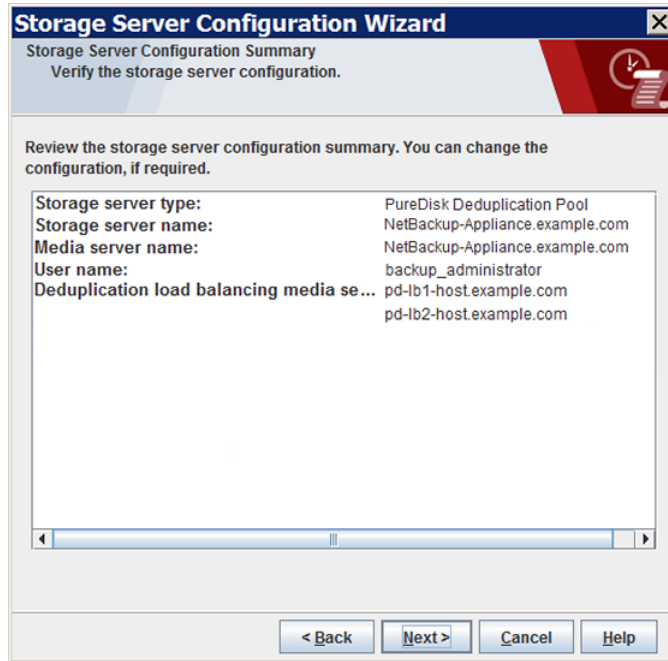
See [“About MSDP load balancing servers”](#) on page 28.

The following is an example of the wizard panel:




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

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



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

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



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

- | | |
|-----------------------|---|
| Configure a disk pool | Ensure that Create a disk pool using the storage server that you have just created is selected and then click Next .

The Select Volumes panel appears. Continue to the next step. |
| Exit | Click Close .

You can configure a disk pool at another time. |

If storage server creation fails, troubleshooting help is available.
See [“Troubleshooting MSDP configuration issues”](#) on page 414.

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

The following is an example of the wizard panel:

Storage Server Configuration Wizard

Volume Selection
Select volumes to use in the disk pool.

Storage server: 'NetBackup-Appliance.example.com'

Storage server type: PureDisk

Select storage server volumes to add to the disk pool.

Volume Name	Available Space	Raw Size	Replication
<input checked="" type="checkbox"/> PureDiskVolume	22.21 GB	22.25 GB	None

Disk Pool Size

Total available space: 22.21 GB

Total raw size: 22.25 GB

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


< Back Next > Cancel Help

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

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

See “[Media Server Deduplication Pool properties](#)” on page 116.

The following is an example of the wizard panel:



Storage Server Configuration Wizard

Additional Disk Pool Information
Provide additional disk pool information.

Storage server: NetBackup-Appliance.example.com

Storage server type: PureDisk

Disk Pool Size

Total available space: 22.21 GB

Total raw size: 22.25 GB

Disk Pool name:

Comments:

High water mark: %

Low water mark: %

Maximum I/O Streams

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


☐ Limit I/O streams: per volume

< Back Next > Cancel Help

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

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

The following is an example of the wizard panel:



Storage Server Configuration Wizard

Disk Pool Configuration Summary
 Verify the disk pool configuration.

Review the disk pool configuration summary. You can change the configuration, if required.


Storage server:	NetBackup-Appliance.example.com
Storage server type:	PureDisk
Volumes:	PureDiskVolume
Disk Pool Details:	
Disk Pool name:	dp-NetBackup-Appliance
High water mark:	98
Low water mark:	80
Maximum IO Streams:	Unlimited
Comments:	

< Back Next > Cancel Help

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

- 11** The **Disk Pool Configuration Status** wizard panel describes the progress of the operation.

The following is an example of the wizard panel:



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

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

Exit Click **Close**.

You can configure one or more storage units later.

See [“Configuring a Media Server Deduplication Pool storage unit”](#) on page 121.

- 12 Enter the appropriate information for the storage unit.

See “Media Server Deduplication Pool storage unit properties” on page 123.

The following is an example of the wizard panel:

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

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

About disk pools for NetBackup deduplication

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

Two types of deduplication pools exist, as follows:

- A NetBackup **Media Server Deduplication Pool** represents the disk storage that is attached to a NetBackup media server. NetBackup deduplicates the data and hosts the storage.

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

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

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

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

Table 5-12 Deduplication pools for duplication or replication

Type	Requirements
Optimized duplication within the same NetBackup domain	<p>Optimized duplication in the same domain requires the following deduplication pools:</p> <ul style="list-style-type: none">■ 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. <p>See “About MSDP optimized duplication within the same domain” on page 141.</p>
Auto Image Replication to a different NetBackup domain	<p>Auto Image Replication deduplication pools can be either replication source or replication target. The replication properties denote the purpose of the deduplication pool. The deduplication pools inherit the replication properties from their volumes.</p> <p>See “About the replication topology for Auto Image Replication” on page 164.</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 159.</p>

Configuring a disk pool for deduplication

The NetBackup **Storage Server Configuration Wizard** lets you configure one disk pool during storage server configuration. To configure additional disk pools, launch the **Disk Pool Configuration Wizard**. Before you can configure a NetBackup disk pool, a NetBackup deduplication storage server must exist.

See [“About disk pools for NetBackup deduplication”](#) on page 107.

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

- The type of disk pool: **PureDisk**. **PureDisk** is the type for the following disk pools:
 - A **Media Server Deduplication Pool** on the disk storage that is attached to a NetBackup deduplication media server.
 - A **PureDisk Deduplication Pool**, which represents storage on a NetBackup 5000 series appliance.
- The deduplication storage server to query for the disk storage to use for the pool.
- The disk volume to include in the pool.
NetBackup exposes the storage as a single volume.
- The disk pool properties.
See [“Media Server Deduplication Pool properties”](#) on page 116.

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

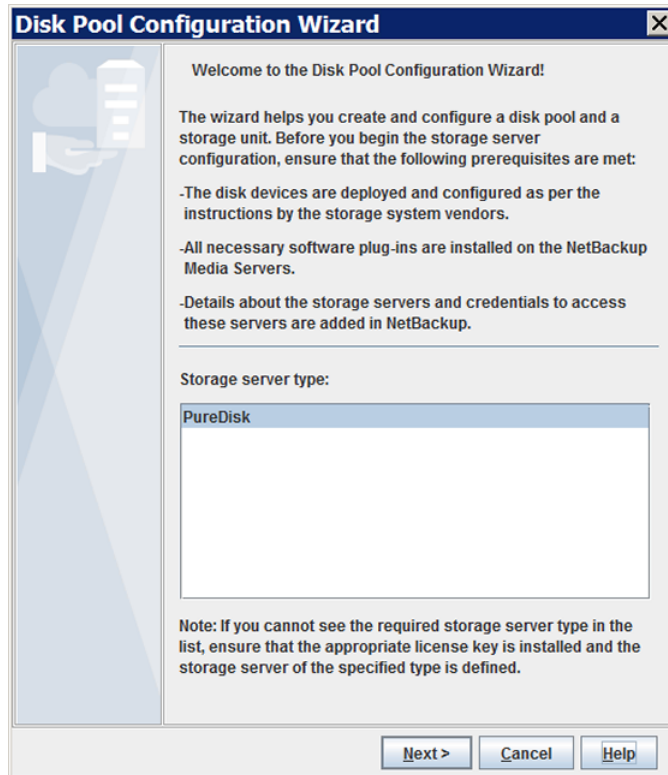
To configure a deduplication disk pool by using the wizard

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

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

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

The following is an example of the wizard panel:



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

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

The following is an example of the wizard panel:

Disk Pool Configuration Wizard

Storage Server Selection
Select storage servers to scan for disk volumes.

Storage server:

Name	Type
msdp-ss-host.example.com	PureDisk

Note: If you cannot see a required storage server in the list, ensure that the storage server details are added in NetBackup.

< Back Next > Cancel Help

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

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

Disk Pool Configuration Wizard

Volume Selection
Select volumes to use in the disk pool.

Storage server: msdp-ss-host.example.com

Storage server type: PureDisk

Select storage server volumes to add to the disk pool.

Volume Name	Available Space	Raw Size	Replication
<input checked="" type="checkbox"/> PureDiskVolume	22.21 GB	22.25 GB	None

Disk Pool Size

Total available space: 22.21 GB

Total raw size: 22.25 GB

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

< Back

Next >

Cancel

Help

- Media Server
Deduplication Pool

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


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

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

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

See “[Media Server Deduplication Pool properties](#)” on page 116.

The following is an example of the wizard panel:



Disk Pool Configuration Wizard

Additional Disk Pool Information
 Provide additional disk pool information.

Storage server: msdp-ss-host.example.com
 Storage server type: PureDisk

Disk Pool Size
 Total available space: 22.21 GB
 Total raw size: 22.25 GB

Disk Pool name:

Comments:

High water mark: %
 Low water mark: %


Maximum I/O Streams
 ⓘ Concurrent read and write jobs affect disk performance.
 Limit I/O streams to prevent disk overload.
☐ Limit I/O streams: per volume

< Back Next > Cancel Help

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

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

The following is an example of the wizard panel:



The screenshot shows a window titled "Disk Pool Configuration Wizard" with a subtitle "Disk Pool Configuration Summary" and the instruction "Verify the disk pool configuration." The main area contains a summary of the configuration with the following details:


Storage server:	msdp-ss-host.example.com
Storage server type:	PureDisk
Volumes:	PureDiskVolume
Disk Pool Details:	
Disk Pool name:	dp-msdp-ss-host
High water mark:	98
Low water mark:	80
Maximum IO Streams:	Unlimited
Comments:	

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

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

- 9 The **Disk Pool Configuration Status** wizard panel describes the progress of the operation.

The following is an example of the wizard panel:



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

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

Exit Click **Close**.

You can configure one or more storage units later.

See [“Configuring a Media Server Deduplication Pool storage unit”](#) on page 121.

- 10 Enter the appropriate information for the storage unit.

See “Media Server Deduplication Pool storage unit properties” on page 123.

The following is an example of the wizard panel:

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

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

Media Server Deduplication Pool properties

Table 5-13 describes the disk pool properties.

Table 5-13 Media server deduplication pool properties

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

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

Property	Description
Storage server type	For a Media Server Deduplication Pool , the storage type is PureDisk .
Disk volumes	<p>For a Media Server Deduplication Pool, all disk storage is exposed as a single volume.</p> <p>PureDiskVolume is a virtual name for the storage that is contained within the directories you specified for the storage path and the database path.</p>
Total available space	The amount of space available in the disk pool.
Total raw size	The total raw size of the storage in the disk pool.
Disk Pool name	The disk pool name. Enter a name that is unique across your enterprise.
Comments	A comment that is associated with the disk pool.
High water mark	<p>The 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.</p> <p>The High water mark includes the space that is committed to other jobs but not already used.</p> <p>The default is 98%.</p>
Low water mark	The Low water mark has no affect on the PureDiskVolume .
Limit I/O streams	<p>Select to limit the number of read and write streams (that is, jobs) for each volume in the disk pool. A job may read backup images or write backup images. By default, there is no limit. If you select this property, also configure the number of streams to allow per volume.</p> <p>When the limit is reached, NetBackup chooses another volume for write operations, if available. If not available, NetBackup queues jobs until a volume is available.</p> <p>Too many streams may degrade performance because of disk thrashing. Disk thrashing is excessive swapping of data between RAM and a hard disk drive. Fewer streams can improve throughput, which may increase the number of jobs that complete in a specific time period.</p>

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

Property	Description
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>

Creating the data directories for 250-TB MSDP support

NetBackup requires that each storage volume contain a directory named `data`. You must create the `data` directories on the second and third volumes that are required for 250-TB support. (NetBackup creates the required `data` directory on the volume that you specify in the **Storage Server Configuration Wizard**.)

- Prerequisite
- The volumes must be formatted with the file systems that NetBackup supports for MSDP and mounted on the storage server. See [“About provisioning the storage for MSDP”](#) on page 47.
 - The storage server must be configured already. See [“Configuring a storage server for a Media Server Deduplication Pool”](#) on page 80.

To create the data directories for 250-TB MSDP support

- ◆ In both the second and third volumes for the **Media Server Deduplication Pool**, create a `data` subdirectory at the volumes's mount points, as follows:

`mount_point/data`

The following is an example of the mount points for the three required storage volumes:

```

/msdp/vol0 <--- Netbackup creates the data directory in this volume
/msdp/vol1 <--- Create a data directory in this volume
/msdp/vol2 <--- Create a data directory in this volume

```

Adding volumes to a 250-TB Media Server Deduplication Pool

When you configure a storage server for a 250-TB **Media Server Deduplication Pool**, you specify the pathname of the first storage volume. Before you can use the **Media Server Deduplication Pool**, you must add the other two volumes to the disk pool.

See [“About provisioning the storage for MSDP”](#) on page 47.

See [“Configuring a storage server for a Media Server Deduplication Pool”](#) on page 80.

To add the other volumes to a 250-TB Media Server Deduplication Pool

- 1 On the MSDP storage server, use the `crcontrol` utility to add the second and third 32-TB volumes to the disk pool. You must include the `data` directories as part of the pathname. The following is the command syntax:

```
/usr/opens/pdde/pdcr/bin/crcontrol --dsaddpartition pathname
```

The following two examples show how to add the pathnames for `/msdp/vol1` and `/msdp/vol2` volumes:

```
# /usr/opens/pdde/pdcr/bin/crcontrol --dsaddpartition /msdp/vol1/data
Partition /msdp/vol1/data was added successfully.
```

```
# /usr/opens/pdde/pdcr/bin/crcontrol --dsaddpartition /msdp/vol2/data
Partition /msdp/vol2/data was added successfully
```

After the next polling cycle of the NetBackup Remote Manager and Monitor Service, the NetBackup Administration Console displays the new, expanded capacity for the disk pool.

Note: For a **Media Server Deduplication Pool**, NetBackup exposes the storage as a single volume in the **NetBackup Administration Console**. Therefore, if you have a 250-TB **Media Server Deduplication Pool**, the number of volumes is 1 even though three volumes are used for the storage.

- 2 On the MSDP storage server, you can use the `crcontrol` utility to verify that the **Media Server Deduplication Pool** contains the new volumes. If the volumes were added correctly, there should be three mount points as follows:

```
/usr/opens/pdde/pdcr/bin/crcontrol --dsstat 2 | grep Mount
```

```
Mount point count: 3
===== Mount point 1 =====
===== Mount point 2 =====
===== Mount point 3 =====
```

- 3 To display detailed information about the disk pool, use the `crcontrol --dsstat 3` command, as follows:

```
/usr/opens/pdde/pdcr/bin/crcontrol --dsstat 3
```

Configuring a Media Server Deduplication Pool storage unit

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

See [“About disk pools for NetBackup deduplication”](#) on page 107.

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 [“Media Server Deduplication Pool storage unit properties”](#) on page 123.

Media Server Deduplication Pool storage unit properties

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

Table 5-14 Media Server Deduplication Pool storage unit properties

Property	Description
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 .
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, Veritas recommends that you use the default, maximum fragment size to ensure optimal deduplication performance. For more information, see the <i>NetBackup Snapshot Client Administrator's Guide</i> : http://www.veritas.com/docs/DOC5332

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

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

MSDP storage unit recommendations

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

Configure a favorable client-to-server ratio

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

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

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

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

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

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

Throttle traffic to the media servers

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

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

Configuring client attributes for MSDP client-side deduplication

To configure client deduplication, set an attribute in the NetBackup master server **Client Attributes** host properties. If the client is in a backup policy in which the

storage destination is a **Media Server Deduplication Pool**, the client deduplicates its own data.

To specify the clients that deduplicate backups

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the details pane, select the master server.
- 3 On the **Actions** menu, select **Properties**.
- 4 Select the **Client Attributes** properties.
- 5 On the **General** tab of the **Client Attributes** properties, add the clients that you want to deduplicate their own data to the **Clients** list, as follows:
 - Click **Add**.
 - In the **Add Client** dialog box, enter a client name or browse to select a client. Then click **Add**.
Repeat for each client that you want to add.
 - When you finish adding clients, click **Close**.
- 6 Select one of the following **Deduplication Location** options:
 - **Always use the media server** disables client deduplication. By default, all clients are configured with the **Always use the media server** option.
 - **Prefer to use client-side deduplication** uses client deduplication if the deduplication plug-in is active on the client. If it is not active, a normal backup occurs; client deduplication does not occur.
 - **Always use client-side deduplication** uses client deduplication. If the deduplication backup job fails, NetBackup retries the job.

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

Note: NetBackup Client Direct deduplication is supported in a Cloud Catalyst duplication and also direct or primary backup to cloud scenario. Upgrade the load balancing media servers for a particular Cloud Catalyst server and the affected clients to NetBackup 8.2.1.

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

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

See “[Disabling MSDP client-side deduplication for a client](#)” on page 127.

Disabling MSDP client-side deduplication for a client

You can remove a client from the list of clients that deduplicate their own data. If you do so, a deduplication server backs up the client and deduplicates the data.

To disable MSDP client deduplication for a client

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the details pane, select the master server.
- 3 On the **Actions** menu, select **Properties**.
- 4 On the **Host Properties Client Attributes General** tab, select the client that deduplicates its own data.
- 5 In the **Deduplication Location** drop-down list, select **Always use the media server**.
- 6 Click **OK**.

About MSDP compression

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

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

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

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

See [“MSDP compression and encryption settings matrix”](#) on page 130.

Table 5-15 MSDP compression options

Option	Description
Compression for backups	<p>For backups, the deduplication plug-in compresses the data after it is deduplicated. The data remains compressed during transfer from the plug-in to the NetBackup Deduplication Engine on the storage server. The Deduplication Engine writes the encrypted data to the storage. For restore jobs, the process functions in the reverse direction.</p> <p>The <code>COMPRESSION</code> parameter in the <code>pd.conf</code> file on each MSDP host controls compression and decompression for that host. By default, backup compression is enabled on all MSDP hosts. Therefore, compression and decompression occur on the following hosts as necessary:</p> <ul style="list-style-type: none"> ■ The clients that deduplicate their own data (that is, client-side deduplication). ■ The load balancing servers. ■ The storage server. <p>MSDP compression cannot occur on normal NetBackup clients (that is, the clients that do not deduplicate their own data).</p> <p>Note: Do not enable backup compression by selecting the Compression option on the Attributes tab of the Policy dialog box. If you do, NetBackup compresses the data before it reaches the plug-in that deduplicates it. Consequently, deduplication rates are very low. Also, NetBackup does not use the Deduplication Multi-Threaded Agent if policy-based encryption is configured.</p> <p>See “About the MSDP Deduplication Multi-Threaded Agent” on page 59.</p>
Compression for duplication and replication	<p>For duplication and replication, the deduplication plug-in compresses the data for transfer. The data remains compressed during transfer from the plug-in to the NetBackup Deduplication Engine on the storage server and remains compressed on the storage.</p> <p>The <code>OPTDUP_COMPRESSION</code> parameter in the <code>pd.conf</code> file controls compression for duplication and replication. By default, duplication and replication compression is enabled on all MSDP hosts. Therefore, duplication and replication compression occurs on the following MSDP servers:</p> <ul style="list-style-type: none"> ■ The load balancing servers. ■ The storage server. <p>Duplication and replication compression does not apply to clients.</p> <p>NetBackup chooses the least busy host to initiate and manage each duplication job and replication job. To ensure that compression occurs for all optimized duplication and replication jobs: do not change the default setting of the <code>OPTDUP_COMPRESSION</code> parameter.</p>

See [“About the MSDP pd.conf configuration file”](#) on page 209.

See [“Use MSDP compression and encryption”](#) on page 44.

About MSDP encryption

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

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

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

See [“MSDP compression and encryption settings matrix”](#) on page 130.

Table 5-16 MSDP encryption options

Option	Description
Backup encryption	<p>For backups, the deduplication plug-in encrypts the data after it is deduplicated. The data remains encrypted during transfer from the plug-in to the NetBackup Deduplication Engine on the storage server. The Deduplication Engine writes the encrypted data to the storage. For restore jobs, the process functions in the reverse direction.</p> <p>The MSDP <code>pd.conf</code> file <code>ENCRYPTION</code> parameter controls backup encryption for individual hosts. By default, backup encryption is disabled on all MSDP hosts. If you want backup encryption, you must enable it on the following MSDP hosts:</p> <ul style="list-style-type: none"> ■ The clients that deduplicate their own data (that is, client-side deduplication). ■ The MSDP load balancing servers. ■ The MSDP storage server. <p>See “Configuring encryption for MSDP backups” on page 132.</p> <p>Note: Do not enable backup encryption by selecting the Encryption option on the Attributes tab of the Policy dialog box. If you do, NetBackup encrypts the data before it reaches the plug-in that deduplicates it. Consequently, deduplication rates are very low. Also, NetBackup does not use the Deduplication Multi-Threaded Agent if policy-based encryption is configured.</p> <p>See “About the MSDP Deduplication Multi-Threaded Agent” on page 59.</p>

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

Option	Description
Duplication and replication encryption	<p>For duplication and replication, the deduplication plug-in on MSDP servers encrypts the data for transfer. The data is encrypted during transfer from the plug-in to the NetBackup Deduplication Engine on the target storage server and remains encrypted on the target storage.</p> <p>The MSDP <code>pd.conf</code> file <code>OPTDUP_ENCRYPTION</code> parameter controls duplication and replication encryption for individual hosts. By default, duplication and replication encryption is disabled on the MSDP storage server and on the MSDP load balancing servers. If you want duplication and replication encryption, you must enable it on the following MSDP servers:</p> <ul style="list-style-type: none"> ■ The load balancing servers. ■ The storage server. <p>Duplication and replication encryption does not apply to clients.</p> <p>NetBackup chooses the least busy host to initiate and manage each duplication job and replication job.</p> <p>See “Configuring encryption for MSDP optimized duplication and replication” on page 133.</p>

NetBackup 8.0 introduced the Advanced Encryption Standard 256 bit, CTR (AES) encryption algorithm to Media Server Deduplication Pool (MSDP). The AES encryption algorithm replaces the older Blowfish encryption algorithm.

See [“About the rolling data conversion mechanism for MSDP”](#) on page 134.

See [“MSDP encryption behavior and compatibilities”](#) on page 137.

MSDP compression and encryption settings matrix

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

Table 5-17 Compression and encryption settings and outcomes

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

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

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

See [“Configuring encryption for MSDP backups”](#) on page 132.

See [“About MSDP compression”](#) on page 127.

See [“About MSDP encryption”](#) on page 129.

See [“About the MSDP `pd.conf` configuration file”](#) on page 209.

Configuring encryption for MSDP backups

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

- | | |
|--|--|
| Configure encryption on individual hosts | <p>Use this procedure to configure encryption on individual MSDP hosts.</p> <p>The <code>ENCRYPTION</code> parameter in the MSDP <code>pd.conf</code> file controls encryption for that host. The parameter applies only to the host on which you modify the <code>pd.conf</code>, as follows:</p> <p>See “To configure backup encryption on a single host” on page 132.</p> |
| Configure encryption for all Client Direct clients | <p>Use this procedure to configure encryption for all of your clients that deduplicate their own data (that is, client-side deduplication). If you use this procedure, you do not have to configure each client-side deduplication client individually.</p> <p>The <code>ServerOptions</code> parameter in the MSDP <code>contentrouter.cfg</code> file controls encryption for all client-side deduplication clients. This parameter supersedes the <code>pd.conf</code> file <code>ENCRYPTION</code> setting on client-side deduplication hosts.</p> <p>See “To configure backup encryption on all client-side deduplication clients” on page 133.</p> |

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

See [“About MSDP encryption”](#) on page 129.

To configure backup encryption on a single host

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

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

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

See [“MSDP pd.conf file parameters”](#) on page 210.

- 2 For the line that begins with `#ENCRYPTION`, remove the pound sign (or hash sign, `#`) in column 1.

- 3 In that same line, replace the 0 (zero) with a 1.

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

- 4 On the client-side deduplication clients and on the MSDP load balancing servers, ensure that the `LOCAL_SETTINGS` parameter in the `pd.conf` file is set to 1. Doing so ensures that the setting on the current host has precedence over the server setting.
- 5 Save and close the file.
- 6 If the host is the storage server or a load balancing server, restart the NetBackup Remote Manager and Monitor Service (`nbrmms`) on the host.

To configure backup encryption on all client-side deduplication clients

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

```
ServerOptions=fast,verify_data_read,encrypt
```

Configuring encryption for MSDP optimized duplication and replication

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

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

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

See [“About MSDP encryption”](#) on page 129.

To configure replication and duplication encryption on an MSDP server

- 1 Use a text editor to open the `pd.conf` file on the MSDP server.
The `pd.conf` file resides in the following directories:
 - (UNIX) `/usr/opensv/lib/ost-plugins/`
 - (Windows) `install_path\Veritas\NetBackup\bin\ost-plugins`
- 2 For the line that begins with `#OPTDUP_ENCRYPTION`, remove the pound sign (or hash sign, #) in column 1.
- 3 In that same line, replace the 0 (zero) with a 1.

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

See [“MSDP pd.conf file parameters”](#) on page 210.

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

About the rolling data conversion mechanism for MSDP

To ensure that data is encrypted and secured with the highest standards, NetBackup uses the AES encryption algorithm and SHA-2 fingerprinting algorithm beginning with the 8.1 release. Specifically, MSDP uses AES-256 and SHA-512/256.

In NetBackup 8.1, with the introduction of the AES and the SHA-2 algorithms, we want to convert the data that is encrypted and computed with the older algorithms (Blowfish and MD5-like) to the newer algorithms (AES-256 and SHA-512/256).

The environments that are upgraded to NetBackup 8.1 may include Blowfish encrypted data and the MD5-like fingerprints that need to be converted to the new format. To handle the conversion and secure the data, a new internal task converts the current data container to the AES-256 encryption and the SHA-512/256 fingerprint algorithm. This new task is referred to as the rolling data conversion.

The conversion begins automatically after an upgrade to NetBackup 8.0. You can control some aspects of the conversion process or stop it entirely.

Rolling data conversion traverses all existing data containers. If the data is encrypted with the Blowfish algorithm, the data is re-encrypted with the AES-256 algorithm. Then a new SHA-512/256 fingerprint is generated. After the conversion, the data container has an additional `.map` file, in addition to the `.bhd` and `.bin` files. The `.map` file contains the mapping between the SHA-512/256 and the MD5-like fingerprints. It is used for the compatibility between SHA-512/256 fingerprints and MD5-like fingerprints. The `.bhd` file includes the SHA-512/256 fingerprints.

When you upgrade to NetBackup 8.1.1, there might be encrypted data that is not encrypted using a customer key. The encrypted data must be encrypted by a customer key and to handle the data conversion and secure the data, a new internal task encrypts the existing data using a customer key. After the encryption and the fingerprint rolling conversion completes, the KMS rolling conversion begins.

Modes of rolling data conversion

MSDP uses the rolling data conversion mechanism to convert Blowfish encrypted data to AES-256 encrypted data, and MD5-like fingerprints to SHA-512/256 fingerprints, in parallel. There are two modes of data conversion: normal mode and fast mode.

- Normal mode: By default, the data conversion process starts in a normal mode for an upgraded system. Similar to compaction, the data conversion runs only when no backup, restore, or Content Router Queue Processing (CRQP) jobs are active.

In the normal mode, the time for data conversion depends on the following factors:

- Total size of the storage
- Power of the CPU
- Workload on the system

Data conversion in the normal mode may take a longer time.

Veritas tests in a controlled environment showed that for a single 1-TB mount point, the conversion speed is about 50MB/s in normal mode.

- Fast mode: In the fast mode, the data conversion disables cyclic redundancy checks and compaction. The rolling data conversion runs while backup, restore, duplication, or CRQP jobs are active.

Veritas tests in a controlled environment showed that for a single 1-TB mount point, the conversion speed is about 105MB/s in fast mode.

Note: The performance numbers shown were observed in the Veritas test environment and are not a guarantee of performance in your environment.

In a new installation of NetBackup 8.1, the rolling data conversion is marked as **Finished** and does not start in the future. For an upgrade to NetBackup 8.1, the rolling data conversion is enabled by default and works in the background after the MSDP conversion completes. Only the data that existed before upgrade is converted. All new data uses the new SHA-512/256 fingerprint and does not need conversion.

While in **Fast mode**, the rolling data conversion affects the performance of backup, restore, duplication, and replication jobs. To minimize this effect, use the **Normal mode**, which pauses the conversion when the system is busy, but slows down the conversion process. The **Fast mode** keeps the conversion active regardless of system state.

You can manage and monitor the rolling data conversion using the following `crcontrol` command options.

Table 5-18 MSDP `crcontrol` command options for rolling data conversion

Option	Description
<code>--dataconverton</code>	<p>To start the data conversion process, use the <code>--dataconverton</code> option:</p> <p>Windows:</p> <pre>install_path\Veritas\pdde\Crcontrol.exe --dataconverton</pre> <p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --dataconverton</code></p>
<code>--dataconvertoff</code>	<p>To stop the data conversion process, use the <code>--dataconverton</code> option:</p> <p>Windows:</p> <pre>install_path\Veritas\pdde\Crcontrol.exe --dataconvertoff</pre> <p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --dataconvertoff</code></p>

Table 5-18 MSDP `crcontrol` command options for rolling data conversion
(continued)

Option	Description
<code>--dataconvertstate</code>	<p>To determine the mode of data conversion and the conversion progress, use the <code>--dataconvertstate</code> option:</p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --dataconvertstate</code></p> <p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --dataconvertstate</code></p>
<code>--dataconvertmode</code>	<p>To switch between the normal mode and fast mode of data conversion, use the <code>--dataconvertmode</code> option:</p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --dataconvertmode mode</code></p> <p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --dataconvertmode <mode></code></p> <p>The default value for the <code><mode></code> variable is 0, which stands for the normal mode. To switch data conversion from normal mode to fast mode, enter 1 for the value of the <code><mode></code> variable.</p>

MSDP encryption behavior and compatibilities

MSDP supports multiple encryption algorithms. Therefore, it manages both the Blowfish and the AES encrypted data to ensure data compatibility.

For restore operations, MSDP recognizes the Blowfish data and the AES data to be able to restore the old backup images.

The following tables describe the encryption behavior for backup, duplication, and replication operations when the encryption is in progress.

Table 5-19 Encryption behavior for a backup operation to a NetBackup 8.0 storage server

Type of client	Data encryption format
Client with NetBackup 8.0, including the Client Direct deduplication	AES

Table 5-19 Encryption behavior for a backup operation to a NetBackup 8.0 storage server (*continued*)

Type of client	Data encryption format
Client with NetBackup version earlier than 8.0, excluding Client Direct deduplication	AES
Client with NetBackup version earlier than 8.0, using the Client Direct deduplication	AES (using inline data conversion)
Load balancing server with NetBackup version 8.0	AES
Load balancing server with NetBackup version earlier than 8.0	AES (using inline data conversion)

Table 5-20 Encryption behavior for optimized duplication and Auto Image Replication operations to a NetBackup 8.0 target server

Type of source storage	Data encryption format for the duplication or the replication data that is encrypted with AES	Data encryption format for the duplication or the replication data that is encrypted with Blowfish
Source server with NetBackup 8.0	AES	AES (using inline data conversion)
Source server with NetBackup version earlier than 8.0	Not applicable	AES (using inline data conversion)

Note: Inline data conversion takes place simultaneously while the backup, duplication, or replication operations are in progress.

Configuring optimized synthetic backups for MSDP

To configure optimized synthetic backups for MSDP, you must select the **Synthetic backup** policy attribute.

To configure optimized synthetic backups for MSDP

- 1 Configure a **Standard** or **MS-Windows** backup policy.
 See “[Creating a backup policy](#)” on page 200.
 See the *NetBackup Administrator's Guide, Volume I*:
<http://www.veritas.com/docs/DOC5332>
- 2 Select the **Synthetic backup** attribute on the **Schedule Attribute** tab of the backup policy.
 See “[Setting MSDP storage server attributes](#)” on page 340.
 See “[Creating a backup policy](#)” on page 200.

About a separate network path for MSDP duplication and replication

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

Backups and restores	<p>For backups and restores, NetBackup uses the network interface that was configured during the storage server configuration.</p> <p>Both the backup and restore traffic and the control traffic travel over the <i>backup</i> network.</p> <p>See “About the network interface for MSDP” on page 35.</p>
Duplication and replication	<p>For the duplication and the replication traffic, configure your host operating systems to use a different network than the one you use for backups and restores.</p> <p>Both the duplication and the replication data traffic and the control traffic travel over the <i>duplication and replication</i> network.</p> <p>See “Configuring a separate network path for MSDP duplication and replication” on page 140.</p> <p>When you configure the optimized duplication or replication target, ensure that you select the host name that represents the <i>duplication and replication</i> network.</p>

- See “[About MSDP optimized duplication within the same domain](#)” on page 141.
- See “[About MSDP replication to a different domain](#)” on page 156.

Configuring a separate network path for MSDP duplication and replication

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

See [“About a separate network path for MSDP duplication and replication”](#) on page 139.

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

The following are the prerequisites:

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

To configure a separate network path for MSDP duplication and replication

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

```
10.10.10.1    TargetStorageServer.example.com    TargetStorageServer
```

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

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

```
10.80.25.66    SourceStorageServer.example.com    SourceStorageServer
```

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

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

```
SourceStorageServer.example.com> ping TargetStorageServer.example.com
TargetStorageServer.example.com> ping SourceStorageServer.example.com
```

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

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

About MSDP optimized duplication within the same domain

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

The optimized duplication operation is more efficient than normal duplication. Only the unique, deduplicated data segments are transferred. Optimized duplication reduces the amount of data that is transmitted over your network.

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

The following table describes the supported sources and targets.

Table 5-21 MSDP optimized duplication source and targets

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

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

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

See “[Configuring NetBackup optimized duplication or replication behavior](#)” on page 153.

You can use a separate network for the duplication traffic.

See “[About a separate network path for MSDP duplication and replication](#)” on page 139.

See “[Configuring MSDP optimized duplication within the same NetBackup domain](#)” on page 149.

Review the following requirements and limitations.

About MSDP optimized duplication requirements

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

- The source storage and the destination storage must have at least one media server in common.
See “[About the media servers for MSDP optimized duplication within the same domain](#)” on page 144.
- In the storage unit you use for the destination for the optimized duplication, you must select only the common media server or media servers.
If you select more than one, NetBackup assigns the duplication job to the least busy media server. If you select a media server or servers that are not in common, the optimized duplication job fails.
For more information about media server load balancing, see the *NetBackup Administrator's Guide, Volume I*:
<http://www.veritas.com/docs/DOC5332>
- The destination storage unit cannot be the same as the source storage unit.

About MSDP optimized duplication limitations

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

- If an optimized duplication job fails after the configured number of retries, NetBackup does not run the job again.
By default, NetBackup retries an optimized duplication job three times. You can change the number of retries.
See “[Configuring NetBackup optimized duplication or replication behavior](#)” on page 153.

- NetBackup does not support MSDP optimized duplication to storage unit groups. If you use a storage unit group as a destination for optimized duplication, NetBackup uses regular duplication.
- Optimized duplication does not support multiple copies. If NetBackup is configured to make multiple new copies from the (source) copy of the backup image, the following occurs:
 - In a storage lifecycle policy, one duplication job creates one optimized duplication copy. If multiple optimized duplication destinations exist, a separate job exists for each destination. This behavior assumes that the device for the optimized duplication destination is compatible with the device on which the source image resides.
If multiple remaining copies are configured to go to devices that are not optimized duplication capable, NetBackup uses normal duplication. One duplication job creates those multiple copies.
 - For other duplication methods, NetBackup uses normal duplication. One duplication job creates all of the copies simultaneously. The other duplication methods include the following: NetBackup Vault, the `bpduplicate` command line, and the duplication option of the **Catalog** utility in the **NetBackup Administration Console**.
- For optimized duplication between two **PureDisk Deduplication Pools**, NetBackup reports that the deduplication rate is 100%. However, deduplication does not occur during optimized duplication. Therefore, ignore the deduplication rate statistics.

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

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

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

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

For optimized **Media Server Deduplication Pool** duplication within the same domain, the source storage and the destination storage must have at least one media server in common. The common server initiates, monitors, and verifies the duplication operation. The common server requires credentials for both the source storage and the destination storage. (For deduplication, the credentials are for the NetBackup Deduplication Engine, not for the host on which it runs.)

Which media server initiates the duplication operation determines if it is a push or a pull operation, as follows:

- If the media server is co-located physically with the source storage server, it is push duplication.
- If the media server is co-located physically with the destination storage server, it is a pull duplication.

Technically, no advantage exists with a push duplication or a pull duplication. However, the media server that initiates the duplication operation also becomes the write host for the new image copies.

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

About MSDP push duplication within the same domain

[Figure 5-1](#) shows a push configuration for optimized duplication within the same domain. The local deduplication node contains normal backups; the remote deduplication node is the destination for the optimized duplication copies. Load balancing server LB_L2 has credentials for both storage servers; it is the common server.

Figure 5-1 Push duplication environment

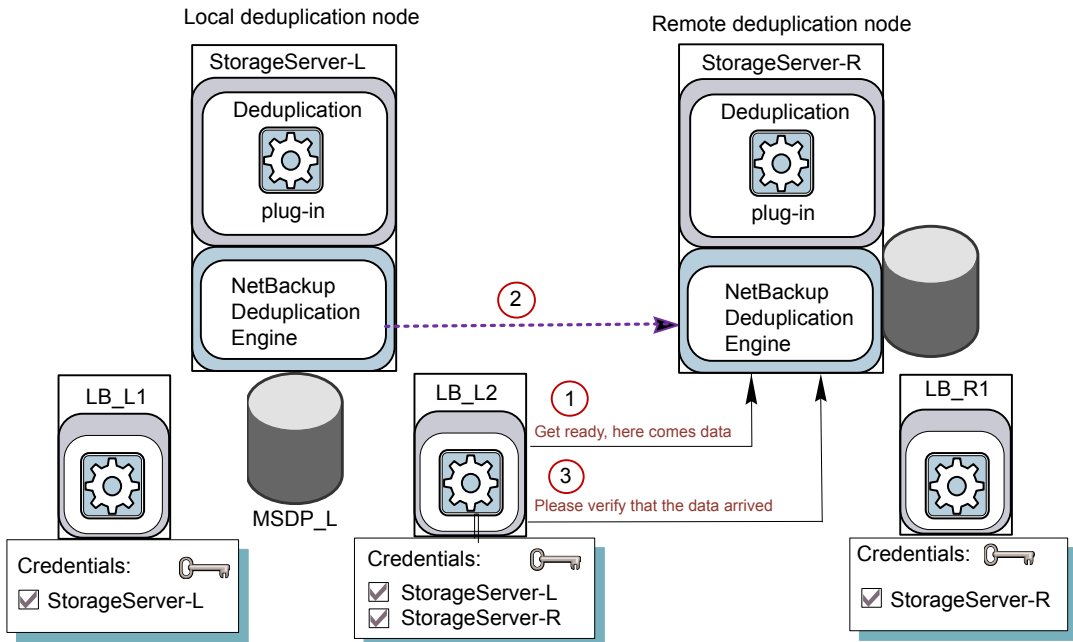


Figure 5-2 shows the settings for the storage unit for the normal backups for the local deduplication node. The disk pool is the **MSDP_L** in the local environment. Because all hosts in the local node are co-located, you can use any available media server for the backups.

Figure 5-2 Storage unit settings for backups to MSDP_L

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

- Storage unit name:** STU-Backups
- Storage unit type:** Disk (selected from a dropdown menu)
- On demand only:** Checked (checkbox)
- Disk type:** PureDisk (selected from a dropdown menu)
- Properties section:**
 - Disk Pool:** MSDP_L (selected from a dropdown menu, with a 'Properties' button next to it)
 - Media Server:**
 - ☒ Use any available media server
 - ☐ Only use the following media servers
 - Media Servers list:**
 - ☐ StorageServer-L
 - ☐ LB_L1
 - ☐ LB_L2

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

Figure 5-3 Storage unit settings for duplication to MSDP_R

New Storage Unit

Storage unit name:
STU-Duplicates

Storage unit type:
Disk ☒ On demand only

Disk type:
PureDisk

Properties

Disk Pool:
MSDP_R

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

Media Servers

- ☐ StorageServer-L
- ☐ LB_L1
- ☒ LB_L2

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

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

About MSDP pull duplication within the same domain

Figure 5-4 shows a pull configuration for optimized duplication within the same domain. Deduplication node A contains normal backups; deduplication node B is the destination for the optimized duplication copies. Host B has credentials for both nodes; it is the common server.

Figure 5-4 Pull duplication




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

Figure 5-5 Pull duplication storage unit settings

New Storage Unit

Storage unit name:

STU-Duplicates

Storage unit type:

Disk

☒ On demand only

Disk type:

PureDisk

Properties

Disk Pool:

MediaServer_DedupePool_B

Properties

Media Server:

☐ Use any available media server

☒ Only use the following media servers

Media Servers

☐ Host_A

☒ Host_B

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

Configuring MSDP optimized duplication within the same NetBackup domain

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

Table 5-22 How to configure optimized duplication of deduplicated data

Step	Action	Description
Step 1	Review optimized duplication	See “About MSDP optimized duplication within the same domain” on page 141.

Table 5-22 How to configure optimized duplication of deduplicated data
(continued)

Step	Action	Description
Step 2	Configure the storage servers	<p>See "Configuring a storage server for a Media Server Deduplication Pool" on page 80.</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 144.</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 disk pool for deduplication" on page 109.</p>
Step 4	Configure the storage unit for backups	<p>In the storage unit for your backups, do the following:</p> <ol style="list-style-type: none"> For the Disk type, select PureDisk. For the Disk pool, select your Media Server Deduplication Pool or PureDisk Deduplication Pool. <p>Note: NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p> <p>If you use a pull configuration, do not select the common media server in the backup storage unit. If you do, NetBackup uses it to deduplicate backup data. (That is, unless you want to use it for a load balancing server for the source deduplication node.)</p> <p>See "Configuring a Media Server Deduplication Pool storage unit" on page 121.</p>

Table 5-22 How to configure optimized duplication of deduplicated data
(continued)

Step	Action	Description
Step 5	Configure the storage unit for duplication	<p>Veritas recommends that you configure a storage unit specifically to be the target for the optimized duplication. Configure the storage unit in the deduplication node that performs your normal backups. Do not configure it in the node that contains the copies.</p> <p>In the storage unit that is the destination for your duplicated images, do the following:</p> <ol style="list-style-type: none"> 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: NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p> <p>Also select Only use the following media servers. Then, select the media server or media servers that are common to both the source storage server and the destination storage server. If you select more than one, NetBackup assigns the duplication job to the least busy media server.</p> <p>If you select only a media server (or servers) that is not common, the optimized duplication job fails.</p> <p>See “Configuring a Media Server Deduplication Pool storage unit” on page 121.</p>
Step 6	Configure optimized duplication bandwidth	<p>Optionally, you can configure the bandwidth for replication.</p> <p>See “About configuring MSDP optimized duplication and replication bandwidth” on page 191.</p>
Step 7	Configure optimized duplication behaviors	<p>Optionally, you can configure the optimized duplication behavior.</p> <p>See “Configuring NetBackup optimized duplication or replication behavior” on page 153.</p> <p>See “About configuring MSDP optimized duplication and replication bandwidth” on page 191.</p>

Table 5-22 How to configure optimized duplication of deduplicated data
(continued)

Step	Action	Description
Step 8	Configure a storage lifecycle policy for the duplication	<p>Configure a storage lifecycle policy only if you want to use one to duplicate images. The storage lifecycle policy manages both the backup jobs and the duplication jobs. Configure the lifecycle policy in the deduplication environment that performs your normal backups. Do not configure it in the environment that contains the copies.</p> <p>When you configure the storage lifecycle policy, do the following:</p> <ul style="list-style-type: none"> ■ 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. <p>Note: NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p> <p>See "About storage lifecycle policies" on page 192.</p> <p>See "Creating a storage lifecycle policy" on page 194.</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 that contains your normal backups. These backups are the primary backup copies. <p>Note: NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p> <p>See "About MSDP backup policy configuration" on page 199.</p> <p>See "Creating a backup policy" on page 200.</p>

Table 5-22 How to configure optimized duplication of deduplicated data
(continued)

Step	Action	Description
Step 10	Configure NetBackup Vault for the duplication	<p>Configure Vault duplication only if you use NetBackup Vault to duplicate the images.</p> <p>Configure Vault in the deduplication environment that performs your normal backups. Do not configure it in the environment that contains the copies.</p> <p>For Vault, you must configure a Vault profile and a Vault policy.</p> <ul style="list-style-type: none"> 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 or a PureDisk 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>Note: NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p>
Step 11	Duplicate by using the <code>bpduplicate</code> command	<p>Use the NetBackup <code>bpduplicate</code> command only if you want to duplicate images manually.</p> <p>Duplicate from a Media Server Deduplication Pool or a PureDisk Deduplication Pool to another Media Server Deduplication Pool or another PureDisk Deduplication Pool in the same domain.</p> <p>Note: NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p> <p>See the <i>NetBackup Commands Reference Guide</i>:</p> <p>http://www.veritas.com/docs/DOC5332</p>

Configuring NetBackup optimized duplication or replication behavior

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

Table 5-23 Optimized duplication behavior

Behavior	Description
Duplication by using NetBackup Vault or the <code>bpduplicate</code> command	<p>If you use NetBackup Vault or the <code>bpduplicate</code> command for duplication, you can configure the following behaviors:</p> <ul style="list-style-type: none">■ Number of optimized duplication attempts. You can change the number of times NetBackup retries an optimized duplication job before it fails the jobs. See “To configure the number of duplication attempts” on page 154.■ Optimized duplication failover. By default, if an optimized duplication job fails, NetBackup does not run the job again. You can configure NetBackup to use normal duplication if an optimized duplication job fails. See “To configure optimized duplication failover” on page 155.
Duplication or replication by using a storage lifecycle policy	<p>If a storage lifecycle policy optimized duplication or replication job fails, NetBackup waits 2 hours and retries the job again. NetBackup repeats the retry behavior until the job succeeds or the source backup image expire.</p> <p>You can change the number of hours for the wait period.</p> <p>See “To configure the storage lifecycle policy wait period” on page 155.</p>

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

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

To configure the number of duplication attempts

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

The file must reside on the master server in the following directory (depending on the operating system):
 - UNIX: `/usr/opensv/netbackup/db/config`

- Windows: `install_path\NetBackup\db\config`.

To configure optimized duplication failover

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

```
RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE
```

See [“Setting NetBackup configuration options by using the command line”](#) on page 155.

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

To configure the storage lifecycle policy wait period

- 1 in the **NetBackup Administration Console**, expand **NetBackup Management** > **Host Properties** > **Master Servers**. Select the master server, and 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

Veritas 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 following NetBackup commands:

For a NetBackup server: `bpsetconfig`

For a NetBackup client: `nbsetconfig`

Configuration options are key and value pairs, as shown in the following examples:

- `CLIENT_READ_TIMEOUT = 300`
- `LOCAL_CACHE = NO`
- `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE`
- `SERVER = server1.example.com`

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

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 the appropriate command. The command depends on the operating system and the NetBackup host type (client or server), as follows:

UNIX On a NetBackup client:

```
/usr/opensv/netbackup/bin/nbsetconfig
```

On a NetBackup server:

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

Windows On a NetBackup client:

```
install_path\NetBackup\bin\nbsetconfig.exe
```

On a NetBackup server:

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

- 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.

You can change existing key and value pairs.

You can add key and value pairs.

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

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

Windows: `Ctrl + Z Enter`

UNIX: `Ctrl + D Enter`

About MSDP replication to a different domain

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

[Table 5-24](#) describes the MSDP replication source and targets that NetBackup supports.

Table 5-24 NetBackup media server deduplication replication targets

Source storage	Target storage
Media Server Deduplication Pool	<p>A Media Server Deduplication Pool, which can be hosted on the following systems:</p> <ul style="list-style-type: none"> ■ A NetBackup media server. ■ A NetBackup 5200 series appliance or NetBackup 5300 series appliance. ■ A Cloud Catalyst storage server.
PureDisk Deduplication Pool	<p>A PureDisk Deduplication Pool on a NetBackup 5000 series appliance. A PureDisk Deduplication Pool is a PureDisk Storage Pool exposed to NetBackup by the NetBackup 5000 series appliance .</p> <p>A PureDisk Deduplication Pool can be the destination of replication from one NetBackup domain only. If you try to replicate from MSDP storage in more than one NetBackup domain, replication does not function correctly.</p> <p>Note: NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p>
Cloud Catalyst Storage Server	<ul style="list-style-type: none"> ■ A disk pool for Cloud Catalyst storage server. ■ A Media Server Deduplication Pool

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

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

See [“Configuring NetBackup optimized duplication or replication behavior”](#) on page 153.

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

You can use a separate network for the duplication traffic.

See [“About a separate network path for MSDP duplication and replication”](#) on page 139.

See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 158.

See [“About MSDP optimized duplication and replication”](#) on page 37.

Configuring MSDP replication to a different NetBackup domain

[Table 5-25](#) describes the tasks that are required to replicate backup images from one **Media Server Deduplication Pool** to another in a different NetBackup domain.

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

See [“About a separate network path for MSDP duplication and replication”](#) on page 139.

Table 5-25 NetBackup MSDP replication configuration tasks

Step	Task	Procedure
Step 1	Learn about MSDP replication	See “About MSDP replication to a different domain” on page 156. See “About NetBackup Auto Image Replication” on page 159.
Step 2	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 167.
Step 3	Add the remote storage server as a replication target	See “Configuring a target for MSDP replication to a remote domain” on page 184. See “Viewing the replication topology for Auto Image Replication” on page 165.

Table 5-25 NetBackup MSDP replication configuration tasks (*continued*)

Step	Task	Procedure
Step 4	Configure a storage lifecycle policy	<p>The following are the options when you configure the SLP operations:</p> <ul style="list-style-type: none"> ■ 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 192.</p> <p>See “About the storage lifecycle policies required for Auto Image Replication” on page 192.</p> <p>See “Creating a storage lifecycle policy” on page 194.</p>
Step 5	Configure replication bandwidth	<p>Optionally, you can configure the bandwidth for replication.</p> <p>See “About configuring MSDP optimized duplication and replication bandwidth” on page 191.</p>

About NetBackup Auto Image Replication

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

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

- One-to-one model
A single production datacenter can back up to a disaster recovery site.
- One-to-many model
A single production datacenter can back up to multiple disaster recovery sites.
See [“One-to-many Auto Image Replication model”](#) on page 161.

- 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.

NetBackup supports Auto Image Replication from a **Media Server Deduplication Pool** in one NetBackup domain to a **Media Server Deduplication Pool** in another domain.

NetBackup supports Auto Image Replication from a disk volume in a **Media Server Deduplication Pool** in one NetBackup domain to a disk volume in a **Media Server Deduplication Pool** in another domain.

Notes about Auto Image Replication

- Although Auto Image Replication is a disaster recovery solution, you cannot restore to clients in the primary (or originating) domain from the target master domain.
- Auto Image Replication does not support synthetic backups or optimized synthetic backups.
- Auto Image Replication does not support spanning volumes in a disk pool. NetBackup fails backup jobs to the disk pools that span volumes if the backup job is in a storage lifecycle policy that also contains a replication operation.
- Auto Image Replication does not support replicating from a storage unit group. That is, the source copy cannot be in a storage unit group.
- The ability to perform Auto Image Replication between different versions of NetBackup does not overrule the basic image compatibility rules. For example, a database backup that was taken in one NetBackup domain can be replicated to a NetBackup domain of an earlier version. However, the older server may not be able to successfully restore from the newer image.
See the following article for information about version compatibility and interoperability:
<http://www.netbackup.com/compatibility>
- 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 5-26 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 192.

Table 5-26 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.

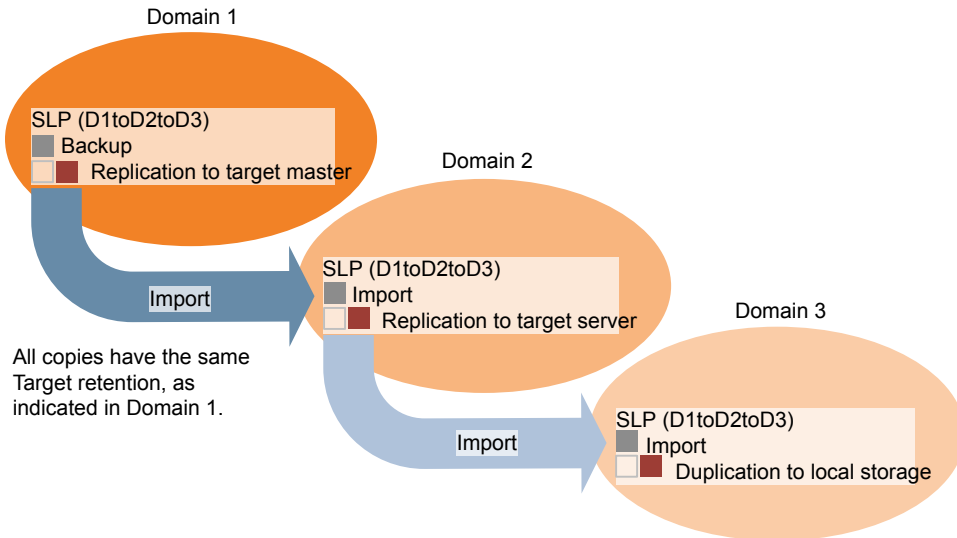
Cascading Auto Image Replication model

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 5-6 represents the following cascading configuration across three domains.

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

Figure 5-6 Cascading Auto Image Replication



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

Note: When the image is replicated in Domain 3, the replication notification event indicates that the master server in Domain 2 is the originating master server. However, after the image is imported successfully into Domain 3, NetBackup correctly indicates that the originating master server is in Domain 1.

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

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

Table 5-27 shows the difference in the import operation setup.

Table 5-27 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.

In the cascading model that is represented in [Figure 5-6](#), all copies have the same **Target Retention**—the **Target Retention** indicated in Domain 1.

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

Figure 5-7 Cascading replications to target master servers, with various target retentions



About the domain relationship for replication

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

See [“Configuring a target for MSDP replication to a remote domain”](#) on page 184.

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 167.

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 184.

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 353.

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 165.

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 164.

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` Use PureDisk for a **Media Server Deduplication Pool**.

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 166.

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.example.com` is the source for the replication operation.
- The `PureDiskVolume` of the deduplication storage server `target_host.dr-site.example.com` is the target of the replication operation.

```
> bpstsinfo -lsuinfo -storage_server bit1.datacenter.example.com -stype PureDisk
```

LSU Info:

```
Server Name: PureDisk:bit1.datacenter.example.com
LSU Name: PureDiskVolume
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: PureDisk storage unit (/bit1.datacenter.example.com#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.example.com: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.example.com -stype PureDisk
LSU Info:
```

```
Server Name: PureDisk:target_host.dr-site.example.com
LSU Name: PureDiskVolume
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: PureDisk storage unit (/target_host.dr-site.example.com#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 for Media Server Deduplication Pool and PureDisk Deduplication Pool as a target storage. Before you configure a Cloud Catalyst storage server as a target storage, establish a trust relationship between the source A.I.R. and the target A.I.R. operations.

The following items describe how a trust relationship affects Auto Image Replication:

No trust relationship

NetBackup replicates to all defined target storage servers.
You cannot select a specific host or hosts as a target.

Trust relationship	You can select a subset of your trusted domains as a target for replication. NetBackup then replicates to the specified domains only rather than to all configured replication targets. This type of Auto Image Replication is known as targeted A.I.R.
--------------------	---

About adding trusted master server using NetBackup CA-signed certificate

With targeted A.I.R., when establishing trust between the source and the remote target server, you need to establish trust in both the domains.

1. In the source master server, add the target master server as a trusted server.
2. In the target master server, add the source master server as a trusted server.

Note: The **NetBackup Administration Console** does not support adding a trusted master server using an external CA-signed certificate.

See [“Adding a trusted master server using external CA-signed certificate”](#) on page 178.

See [“About the certificate to be used for adding a trusted master server”](#) on page 171.

The following diagram illustrates the different tasks for adding trusted master servers when NetBackup CA-signed certificate (or host ID-based certificate) is used for establishing trust between the source and the target master servers.

Figure 5-8 Tasks to establish a trust relationship between master servers for targeted A.I.R. using NetBackup CA-signed certificate

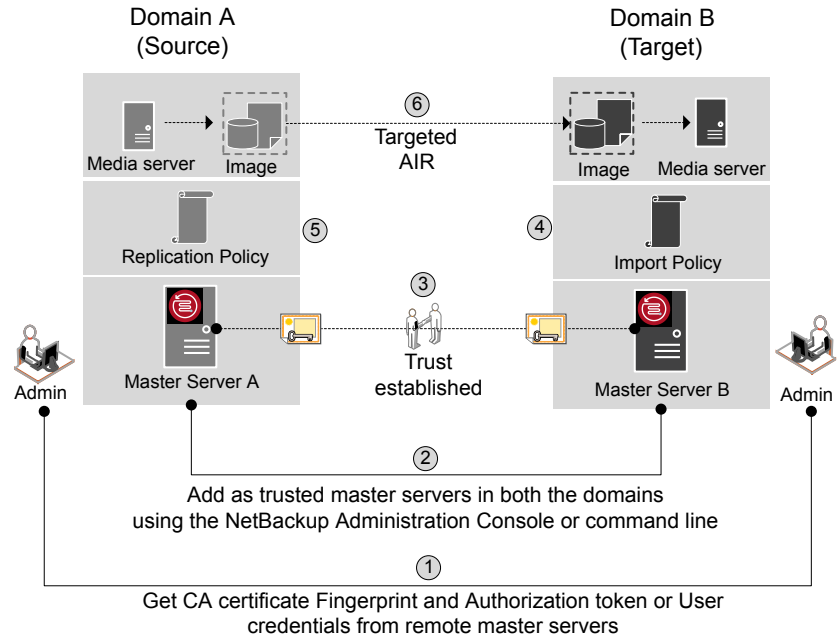


Table 5-28 Tasks to establish a trust relationship between master servers for targeted A.I.R.

Step	Task	Procedure
Step 1	Administrators of both the source and target master servers must obtain each other's CA certificate fingerprint and authorization tokens or the user credentials. This activity must be performed offline. Note: It is recommended to use an authentication token to connect to the remote master server. An authentication token provides restricted access and allows secure communication between both the hosts. The use of user credentials (user name and password) may present a possible security breach.	To obtain the authorization tokens, use the <code>bpnbat</code> command to login and <code>nbcertcmd</code> to get the authorization tokens. To obtain the SHA1 fingerprint of root certificate, use the <code>nbcertcmd -displayCACertDetail</code> command. To perform this task, see the NetBackup Commands Reference Guide . Note: When you run the commands, keep the target as the remote server.

Table 5-28 Tasks to establish a trust relationship between master servers for targeted A.I.R. (*continued*)

Step	Task	Procedure
Step 2	<p>Establish trust between the source and the target domains.</p> <ul style="list-style-type: none"> ■ On the source master server, add the target master server as trusted server. ■ On the target master server, add the source master server as trusted server. 	<p>To perform this task in the NetBackup Administration Console, see the following topic:</p> <p>See “Adding a trusted master server using a NetBackup CA-signed (host ID-based) certificate” on page 173.</p> <p>To perform this task using the <code>nbseccmd</code>, see the NetBackup Commands Reference Guide.</p>
Step 3	<p>After you have added the source and target trusted servers, they have each other's host ID-based certificates. The certificates are used during each communication.</p> <p>Master Server A has a certificate issued by Master Server B and vice versa. Before communication can occur, Master Server A presents the certificate issued by Master Server B and vice versa. The communication between the source and target master servers is now secured.</p>	<p>To understand the use of host ID-based certificates, see the NetBackup Security and Encryption Guide.</p>
Step 3.1	<p>Configure the source media server to get the security certificates and the host ID certificates from the target master server.</p>	<p>See “Configuring NetBackup CA and NetBackup host ID-based certificate for secure communication between the source and the target MSDP storage servers” on page 182.</p> <p>See “Configuring a NetBackup Deduplication Engine user with limited permissions for Auto Image Replication” on page 190.</p>
Step 4	<p>Create an import storage lifecycle policy in the target domain.</p>	<p>See “About storage lifecycle policies” on page 192.</p>
Step 5	<p>On the source MSDP server, use the Replication tab from the Change Storage Server dialog box to add the credentials of the target storage server.</p>	<p>See “Configuring a target for MSDP replication to a remote domain” on page 184.</p>
Step 5.1	<p>Create a replication storage lifecycle policy in the source domain using the specific target master server and storage lifecycle policy.</p> <p>The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains.</p>	<p>See “About storage lifecycle policies” on page 192.</p>

Table 5-28 Tasks to establish a trust relationship between master servers for targeted A.I.R. *(continued)*

Step	Task	Procedure
Step 6	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.	See “About NetBackup Auto Image Replication” on page 159.

If your source and target trusted servers use different NetBackup versions, consider the following.

Note: When you upgrade both the source and the target master server to version 8.1 or later, you need to update the trust relationship. Run the following command:

```
nbseccmd -setuptrustedmaster -update
```

See the [NetBackup Commands Reference Guide](#).

Table 5-29 Trust setup methods for different NetBackup versions

Source server version	Target server version	Trust setup method
8.1 and later	8.1 and later	Add a trusted master server using authorization token. Complete action on both the servers.
8.1 and later	8.0 or earlier	On the source server, add the target as the trusted master server using the remote (target) server's credentials.
8.0 or earlier	8.1 and later	On the source server, add the target as the trusted master server using the remote (target) server's credentials.

About the certificate to be used for adding a trusted master server

Source or target master servers may use NetBackup CA-signed certificates (host ID-based certificates) or external CA-signed certificates.

For more information on NetBackup host ID-based certificates and external CA support, refer to the [NetBackup Security and Encryption Guide](#).

To establish trust between source and target master servers, NetBackup verifies the following:

- Can the source master server establish trust using external CA-signed certificate?

If the external CA configuration options - `ECA_CERT_PATH`, `ECA_PRIVATE_KEY_PATH`, and `ECA_TRUST_STORE_PATH` - are defined in the NetBackup configuration file of the source master server, it can establish the trust using an external certificate.

In case of Windows certificate trust store, only `ECA_CERT_PATH` is defined.

For more information on the configuration options, refer to the [NetBackup Administrator's Guide, Volume I](#).
- Which certificate authorities (CA) does the target master server support?

The target master server may support external CA, NetBackup CA, or both. The following settings show the CA usage information of the master server:

 - In the NetBackup Administration Console - **NetBackup Management > Security Management > Global Security Settings**
 - In the NetBackup web user interface - **Security > Global Security Settings > Secure Communication**.

The following table lists CA support scenarios and certificate to be used to establish trust between the source and the target master servers.

Table 5-30 Certificate to be used for trust setup

Source master server capability to use external certificate	CA usage of the target master server	Certificate to be used for trust setup
Yes The source master server can use NetBackup CA and external CA for communication with a remote master server	External CA	External CA See “Adding a trusted master server using external CA-signed certificate” on page 178.
	NetBackup CA	NetBackup CA See “Adding a trusted master server using a NetBackup CA-signed (host ID-based) certificate” on page 173.
	External CA and NetBackup CA	NetBackup prompts to select the CA that you want to use for trust setup <ul style="list-style-type: none"> ■ If you choose to use external CA, do the following: See “Adding a trusted master server using external CA-signed certificate” on page 178. ■ If you choose to use NetBackup CA, do the following: See “Adding a trusted master server using a NetBackup CA-signed (host ID-based) certificate” on page 173.

Table 5-30 Certificate to be used for trust setup (*continued*)

Source master server capability to use external certificate	CA usage of the target master server	Certificate to be used for trust setup
No	External CA	No trust is established
The source master server can use only NetBackup CA for communication with a remote maser server	NetBackup CA	NetBackup CA See “Adding a trusted master server using a NetBackup CA-signed (host ID-based) certificate” on page 173.
	External CA and NetBackup CA	NetBackup CA See “Adding a trusted master server using a NetBackup CA-signed (host ID-based) certificate” on page 173.

Adding a trusted master server using a NetBackup CA-signed (host ID-based) certificate

Replication operations require that a trust relationship exists between the NetBackup servers in the different domains.

Before you begin

Perform the following steps on both the source and the target server:

- Identify the NetBackup versions that are installed on the source and the target servers.
- Obtain the authorization tokens of the remote server.
Use the `bpnbat` command to log on and `nbcertcmd` to get the authorization tokens.
- Obtain the fingerprints for the remote server.
To obtain the SHA1 fingerprint of root certificate, use the `nbcertcmd -displayCACertDetail` command.
- Ensure that you have one of the following permissions:
 - System administrator permissions with `root` permissions for UNIX, administrator permissions for Windows, or a NetBackupCLI user for a 3.1 NetBackup appliance.
 - Access to the NetBackup Administration Console, where you have `<username> ADMIN=ALL` permissions through `auth.conf`.
 - Enhanced Auditing (EA) user permissions through `authalias.conf`.

- For remote Windows master server, if the user's domain is not same as that of the authentication service, you must add the domain with LDAP using the `vssat addldapdomain` command. See the [NetBackup Commands Reference Guide](#).
Also, this user must have RBAC security administrator permissions. See the [NetBackup Web UI Security Administrator's Guide](#).


Adding a trusted master server, when both the source and the target servers are NetBackup version 8.1 or later

Use this procedure to add a trusted master server when both the source and target servers are NetBackup version 8.1 or later.

See [“Adding a trusted master server using external CA-signed certificate”](#) on page 178.

To add a trusted master server, when both the source and the target servers are NetBackup version 8.1 or later

- 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 and **Actions** > **Properties**.
- 3 In the properties dialog box left pane, select **Servers**.
- 4 On the **Trusted Master Servers** tab, click **Add**.
- 5 Enter the fully-qualified host name of the remote master server and click **Validate Certificate Authority**.



- 6** In the **Validate Certificate Authority** dialog box, verify if the CA certificate fingerprint of the remote server is correct.

To proceed, click **Yes**.

If the fingerprints don't match, click **No**. Contact the remote server admin to provide the correct fingerprints.

- 7** Enter the trusted master server details using one of the following methods.
 - (Recommended) Select **Specify authentication token of the trusted master server** and enter the token details of the remote master server.
 - Select **Specify credentials of the trusted master server** and enter the user name and password. Note that this method may present a possible security breach. Only an authentication token can provide restricted access and allow secure communication between both the hosts.
To establish trust with a 3.1 NetBackup master appliance, use the NetBackup CLI credentials.

Add Trusted Master Server

Establish a trust relationship between your master server and a remote master server to enable communication between the two servers.

Trusted Master Server:

☒ Specify authentication token of the trusted master server

Token:

☐ Show Token

☐ Specify credentials of the trusted master server

User name:

For a Windows host, specify the domain name along with the user name. For example, domain name\user name.

Password:


- 8 Click **OK**.
- 9 Perform the same procedure on the remote master server that you added in step 5.

Adding a trusted master server, when both the source and the target server are NetBackup version 8.0

Use this procedure to add a trusted master server when both the source and target servers are NetBackup version 8.0.

To add a trusted master server, when both the source and the target server are NetBackup version 8.0

- 1 Ensure that the **Enable insecure communication with NetBackup 8.0 and earlier hosts** option is enabled in the global security settings.
- 2 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers** in the left pane.
- 3 In the right pane, select the master server and **Actions > Properties**.
- 4 In the properties dialog box left pane, select **Servers**.
- 5 On the **Trusted Master Servers** tab, click **Add**.
- 6 Enter the fully-qualified host name of the remote master server and click **Validate Certificate Authority**.



- 7 Enter the **Username** and **Password** of the remote master server host.

- 8 Click **OK**.

More information

See [“About trusted master servers for Auto Image Replication”](#) on page 167.

See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 158.

For details on usage reporting in the web UI, see the *NetBackup Web UI for Backup Administrator's Guide*.

For more information on commands, see the [NetBackup Commands Reference Guide](#). For details on the `authalias.conf`, see the [NetBackup Security and Encryption Guide](#).

Adding a trusted master server using external CA-signed certificate

You can now establish a trust between source and target master servers using an external CA-signed certificate.

For more information on the external CA support, refer to the *NetBackup Security and Encryption Guide*.

See [“About the certificate to be used for adding a trusted master server”](#) on page 171.

Note: The **NetBackup Administration Console** does not support adding a trusted master server using an external certificate.

If you try to add a trusted master server with an external certificate using the **NetBackup Administration Console**, an error is displayed.

To add a trusted master server using an external certificate

- 1 Configure the following external certificate configuration options on the source master server:
 - ECA_CERT_PATH
 - ECA_PRIVATE_KEY_PATH
 - ECA_TRUST_STORE_PATH
 - ECA_KEY_PASSPHRASEFILE (optional)

Note: In case of Windows certificate store, configure only the `ECA_CERT_PATH` configuration option.

- 2 Run the `nbseccmd -setuptrustedmaster` command on the source master server.

For more information on the commands, refer to the [NetBackup Commands Reference Guide](#).

If the source and target master servers are configured with external certificates that are issued by different certificate authorities, refer to the following section from the *NetBackup Deduplication Guide*: [Configuring external CA for secure communication between the source MSDP storage server and the target MSDP storage server](#).

Removing a trusted master server

To remove a trusted master server, you must perform the following procedure on both the source and the target server.

Note: If either your source or the target server is on version 8.0 or earlier, follow the procedure that is prescribed in the respective guide.

To remove a trusted master server

- 1 Ensure that all replication jobs to the trusted target master server are complete. You can use `nbstlutil stlilist` to list the state of all storage lifecycle policy-managed operations. To cancel jobs use `nbstlutil cancel`.

See the [NetBackup Commands Reference Guide](#) for information about the `nbstlutil` command.

- 2 Delete all storage lifecycle policies (SLPs) that use the trusted master as a destination.

Note: Before deleting a storage lifecycle policy, ensure that there are no backup policies that indicate the SLP for the **Policy storage**.

- 3 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers** in the left pane.
- 4 In the right pane, select the master server.
- 5 On the **Actions** menu, click **Properties**.
- 6 In the properties dialog box left pane, select **Servers**.
- 7 In the **Servers** dialog box, select the **Trusted Master Servers** tab.
- 8 On the **Trusted Master Servers** tab, select the trusted master server that you want to remove and click **Remove**.

The **Remove Server** confirmation dialog box is displayed.

- 9 Click **Yes**.
- 10 When you finish removing trusted master servers, click **OK**.
- 11 Restart the `nbsl` service.
- 12 Repeat the steps on the source master server.

Note: In case of multiple NICs, if you have established trust using more than one host NIC and if you remove the trust relationship with any one host NIC, the trust with all the other host NICs is broken.

Enabling NetBackup clustered master server inter-node authentication

NetBackup requires inter-node authentication among the master servers in a cluster. For authentication, you must provision an authentication certificate on all of the

nodes of the cluster. The certificates are used to establish SSL connections between the NetBackup hosts. The inter-node authentication allows the following NetBackup functionality:

NetBackup Administration Console	The NetBackup Administration Console in master server clusters requires the NetBackup authentication certificates for correct functionality.
Targeted A.I.R. (Auto Image Replication)	<p>Auto Image Replication in which a master server is in a cluster requires inter-node authentication among the hosts in that cluster. The NetBackup authentication certificates provide the means to establish the proper trust relationships.</p> <p>Provision the certificates on the cluster hosts before you add the trusted master server. This requirement applies regardless of whether the clustered master server is the source of the replication operation or the target.</p> <p>See "About trusted master servers for Auto Image Replication" on page 167.</p>

To enable clustered master server inter-node authentication

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

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

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

The following is example output:

```
# bpnbaz -setupat
You will have to restart Netbackup services on this machine after
the command completes successfully.
Do you want to continue(y/n)y
Gathering configuration information.
Please be patient as we wait for 10 sec for the security services
to start their operation.
Generating identity for host 'bitl.remote.example.com'
Setting up security on target host: bitl.remote.example.com
nbatd is successfully configured on Netbackup Master Server.
Operation completed successfully.
```

Configuring NetBackup CA and NetBackup host ID-based certificate for secure communication between the source and the target MSDP storage servers

MSDP now supports secure communications between two media servers from two different NetBackup domains. The secure communication is set up when you run Auto Image Replication (A.I.R.). The two media servers must use the same CA to do the certificate security check. The source MSDP server uses the CA of the target NetBackup domain and the certificate that is authorized by the target NetBackup domain. You must manually deploy CA and the certificate on the source MSDP server before using Auto Image Replication.

Note: After you upgrade to NetBackup 8.1.2 or later, manually deploy NetBackup CA and the NetBackup host ID-based certificate on the source MSDP server to use the existing Auto Image Replication.

To configure the NetBackup CA and a NetBackup host ID-based certificate, complete the following steps:

1. On the target NetBackup master server, run the following command to display the NetBackup CA fingerprint:

- Windows

```
install_path\NetBackup\bin\NBCertCmd -displayCACertDetail
```

- UNIX

```
/usr/opensv/netbackup/bin/nbcertcmd -displayCACertDetail
```

2. On the source MSDP storage server, run the following command to get the NetBackup CA from target NetBackup master server:

- Windows

```
install_path\NetBackup\bin\NBCertCmd -getCACertificate -server  
target_master_server
```

- UNIX

```
/usr/opensv/netbackup/bin/nbcertcmd -getCACertificate -server  
target_master_server
```

When you accept the CA, ensure that the CA fingerprint is the same as displayed in the previous step.

3. On the source MSDP storage server, run the following command to get a certificate generated by target NetBackup master server:

- Windows

```
install_path\NetBackup\bin\NBCertCmd -getCertificate -server
target_master_server -token token_string
```

- **UNIX**

```
/usr/openv/netbackup/bin/nbcertcmd -getCertificate -server
target_master_server -token token_string
```

4. Use either of these two methods to obtain the authorization tokens:

- **NetBackup Administration Console**

- Log on the target NetBackup master server and open **Security Management > Certificate Management > Token Management**.
- Click the **Create Token** option to create a token, or right-click the blank area of the **Token records** list view and select the **New Token** menu item to create a token.

- **NetBackup Commands**

- Use the `bpbnat` command to log on the target NetBackup master server.
- Use the `nbcertcmd` command to get the authorization tokens.

For more information on the commands, refer to the *NetBackup Commands Reference Guide*.

Configuring external CA for secure communication between the source MSDP storage server and the target MSDP storage server

MSDP now supports use of an external CA for secure communication between two media servers that are from two different NetBackup domains. The secure communication is set up when you run Auto Image Replication (A.I.R.). If the two media servers use different external CAs, then you must exchange the external certificates before you use Auto Image Replication.

To exchange the external certificates, complete the following steps:

1. Copy the root certificate file from the source MSDP storage server to the target MSDP storage server. Combine the certificate files on the target MSDP storage server.
2. Copy the root certificate file from the target MSDP storage server to the source MSDP storage server. Combine the certificate files on the source MSDP storage server.

If the Windows certificate store is used to store the root certificate, add the root certificate to the certificate store. You can use the `certutil` tool to add the root certificate to the certificate store, or just right-click the root certificate file and select

Install Certificate. When you use the `certutil` tool to install the root certificate, the store name parameter must be **Root**. When you use Windows explorer to install the root certificate, the store location must be **Local Machine** and store name must be **Trusted Root Certification Authorities**.

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 156.

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 158.

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 167.


See [“Enabling NetBackup clustered master server inter-node authentication”](#) on page 180.

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

To configure a Media Server Deduplication Pool as a replication target

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

- 4 In the **Change Storage Server** dialog box, select the **Replication** tab.
 The following is an example of the **Change Storage Server** dialog box **Replication** tab:



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

The following is an example of the dialog box.

Add a Replication Target in a Different NetBackup...

Replication Target across Different Domain

Select a trusted master server which manages the target storage device.
 Note: Specifying a target master server is not mandatory if you always want to replicate backups to all targets. Add a target master server only if you want to replicate backups to specific master server domains.

Target master server: <Select Target Master Server from the List>

Target storage server type: PureDisk

Target storage server name:

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

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

User name:

Password:

OK Cancel Help

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

Add a trusted master server

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

To add a trusted master server

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

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

See ["Target options for MSDP replication"](#) on page 188.

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 167.

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 188.

- Enter a replication target

Enter a replication target, except the Cloud Catalyst storage server, 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 188.
- 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 353.

Target options for MSDP replication

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

Table 5-31 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>
Target storage server type	<p>If a trusted master server is configured, the value is Target storage server name.</p> <p>If a trusted master server is not configured, the value is PureDisk.</p>

Table 5-31 MSDP target replication options (*continued*)

Option	Description
Target storage server name	<p>Select the Cloud Catalyst storage server name only from the drop-down list of the trusted master server.</p> <p>If a trusted master server is configured, select the target storage server. If a trusted master server is <i>not</i> configured, enter the name of the target storage server.</p> <p>The drop-down list shows all the storage servers that match the Target storage server type.</p>
User name	<p>When you configure a replication target, NetBackup populates the User name field with user account of the target storage server, as follows:</p> <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. <p>NetBackup supports PureDisk Deduplication Pool storage on NetBackup 5000 series appliances only.</p> <ul style="list-style-type: none"> ■ For a Cloud Catalyst storage server target, the Access Key ID of the cloud account. <p>For additional security, you can give limited permissions to the deduplication engine user.</p> <p>See “Configuring a NetBackup Deduplication Engine user with limited permissions for Auto Image Replication” on page 190.</p> <p>Note: For Cloud Catalyst storage server configured using AWS IAM Role or credential broker type (CREDS_CAP), provide the credentials you had entered when configuring the Cloud Catalyst storage server.</p> <p>See “Managing Cloud Catalyst storage server with IAM Role or CREDS_CAPS credential broker type” on page 270.</p>

Table 5-31 MSDP target replication options (continued)

Option	Description
Password	<p>Enter the password for the NetBackup Deduplication Engine, or enter the secret access key for the Cloud Catalyst storage server.</p> <p>Note: For Cloud Catalyst storage server configured using AWS IAM Role or credential broker type (CREDS_CAP), provide the credentials you had entered when configuring the Cloud Catalyst storage server.</p> <p>See “Managing Cloud Catalyst storage server with IAM Role or CREDS_CAPS credential broker type” on page 270.</p>

Configuring a NetBackup Deduplication Engine user with limited permissions for Auto Image Replication

MSDP supports the creation of a user specifically for Auto Image Replication. A user with permissions limited to Auto Image Replication is more secure than a user with administrative permissions.

To configure a NetBackup Deduplication Engine user with limited permissions for Auto Image Replication, complete the following steps:

- Run the following command on the target MSDP server to add a user for AIR:

Windows

```
<install_path>\pdde\spausers -a -u <username> -p <password> --role air
```

UNIX

```
/usr/opensv/pdde/pdcr/bin/spausers -a -u <username> -p <password> --role air
```
- During configuration of MSDP as a replication target on the source NetBackup master server, enter the user name and password of the user with limited permissions for AIR.

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 209.

See [“Editing the MSDP pd.conf file”](#) on page 210.

See [“MSDP pd.conf file parameters”](#) on page 210.

See “Configuring MSDP optimized duplication within the same NetBackup domain” on page 149.

See “Configuring MSDP replication to a different NetBackup domain” on page 158.

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.veritas.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:

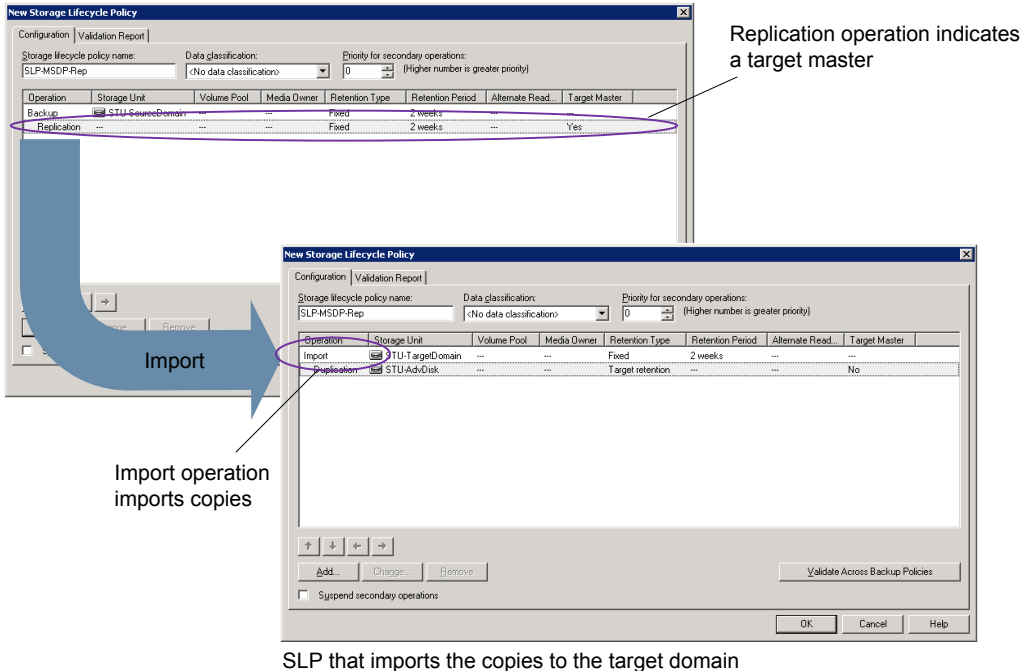
Table 5-32 SLP requirements for Auto Image Replication

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 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 167.■ The SLP must be of the same data classification as the Import SLP in Domain 2.
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 5-9](#) shows how the SLP in the target domain is set up to replicate the images from the originating master server domain.

Figure 5-9 Storage lifecycle policy pair required for Auto Image Replication

SLP on master server in the source domain




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

Creating a storage lifecycle policy

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

To add a storage operation to a storage lifecycle policy

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



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


See [“NetBackup naming conventions”](#) on page 18.

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


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

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

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



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



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

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

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

See [“About NetBackup Auto Image Replication”](#) on page 159.

- 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 200.

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 5-10 Storage Lifecycle Policy tab

Change Storage Lifecycle Policy

Storage Lifecycle Policy | Validation Report

Storage lifecycle policy name: SLP_1_snapshot Data classification: <No data classification> Priority for secondary operations: 0 (Higher number is greater priority)

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

↑ ↓ ← →

Add... Change... Remove

State of Secondary Operation Processing

☒ Active
☐ Postponed
☐ Urgent 12/31/2038 11:59:59 PM

To find impact on Policies associated with this Storage Lifecycle Policy due to change in configuration click here.

Validate Across Backup Policies

OK Cancel Help

Table 5-33 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 192.</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 5-33 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	<p>Use the arrows to indicate the indentation (or hierarchy) of the source for each copy. One copy can be the source for many other copies.</p>
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> <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

Use the following procedure to create a backup policy.

To create a policy

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Policies**.
- 2 Select **Actions > New > Policy**.
- 3 Type a unique name for the policy.
See [“NetBackup naming conventions”](#) on page 18.
- 4 Clear the **Use Policy Configuration Wizard** and click **OK**.
- 5 Configure the attributes, the schedules, the clients, and the backup selections for the new policy.

Resilient Network properties

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

The **Resilient Network** properties let you configure NetBackup to use resilient network connections for backups and restores. A resilient connection allows backup and restore traffic between a client and a NetBackup media server to function effectively in high-latency, low-bandwidth networks such as WANs. The data travels across a wide area network (WAN) to media servers in a central datacenter.

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

The NetBackup Remote Network Transport Service manages the connection between the computers. The Remote Network Transport Service runs on the master

server, the client, and the media server that processes the backup or restore job. If the connection is interrupted or fails, the services attempt to re-establish a connection and synchronize the data.

NetBackup protects only the network socket connections that the NetBackup Remote Network Transport Service (`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 Technology (GRT) for Exchange Server or SharePoint Server
- 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.veritas.org` is a different subdomain than `china.veritas.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 5-11 Master server Resilient Network host properties




Table 5-34 describes the **Resilient Network** properties.

Table 5-34 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, it is recommended 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. It is recommended 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, it is recommended 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 Logging Reference 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 200.

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

- Windows: `install_path\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, it is recommended that you use the fully qualified domain name.

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

- 3 Ensure that the **Resiliency On** option is selected.

- 4 Click **Add**.

- 5 Repeat until you have finished entering clients or address ranges.

- 6 When you finish adding network settings, click **Close**.

Add To All

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

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

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

Change

To change resilient settings

- 1 Select the client host name, the IP address, or the address range.

- 2 Click **Change**.

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

- 3 Select the desired **Resiliency** setting.

- 4 Click **OK**.

Remove

Remove the select host or address rang

- 1 Select the client host name, the IP address, or the address range.

- 2 Click **Remove**.

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



Move an item in the list of items

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

The order of the items in the list is significant.

See [“Resilient Network properties”](#) on page 200.

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

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

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

Adding an MSDP load balancing server


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

See [“About MSDP storage servers”](#) on page 27.

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 variable-length deduplication on NetBackup clients

Currently, NetBackup deduplication follows a fixed-length deduplication method where the data streams are chunked into fixed-length segments (128 KB) and then processed for deduplication. Fixed-length deduplication has the advantage of being a swift method and it consumes less computing resources. Fixed-length deduplication handles most kinds of data streams efficiently. However, there can be cases where fixed-length deduplication might result in low deduplication ratios.

If your data was modified in a shifting mode, that is, if some data was inserted in the middle of a file, then variable-length deduplication enables you to get higher deduplication ratios when you back up the data. Variable-length deduplication reduces backup storage, improves the backup performance, and lowers the overall cost that is spent on data protection.

Note: Use variable-length deduplication for data types that do not show a good deduplication ratio with the current MSDP intelligent deduplication algorithm and affiliated streamers. Enabling Variable-length deduplication might improve the deduplication ratio, but consider that the CPU performance might get affected.

In variable-length deduplication, every segment has a variable size with configurable size boundaries. The NetBackup client examines and applies a secure hash algorithm (SHA-2) to the variable-length segments of the data. Each data segment is assigned a unique ID and NetBackup evaluates if any data segment with the same ID exists in the backup. If the data segment already exists, then the segment data is not stored again.

Warning: If you enable compression for the backup policy, variable-length deduplication does not work even when you configure it.

The following table describes the effect of variable-length deduplication on the data backup:

Table 5-35 Effect of variable-length deduplication

Effect on the deduplication ratio	Variable-length deduplication is beneficial if the data file is modified in a shifting mode, that is when data is inserted, removed, or modified at a binary level. When such modified data is backed up again, variable-length deduplication achieves a higher deduplication ratio. Thus, the second or subsequent backups have higher deduplication ratios.
Effect on the CPU	Variable-length deduplication can be a bit more resource-intensive than fixed-length deduplication to achieve a better deduplication ratio. Variable-length deduplication needs more CPU cycles to compute segment boundaries and the backup time might be more than the fixed-length deduplication method.

Table 5-35 Effect of variable-length deduplication (*continued*)

Effect on data restore	Variable-length deduplication does not affect the data restore process.
------------------------	---

Configure variable-length deduplication

By default, the variable-length deduplication is disabled on a NetBackup client. You can enable variable-length deduplication by adding parameters in the `pd.conf` file. To enable the same settings for all NetBackup clients or policies, you must specify all the clients or policies in the `pd.conf` file.

In a deduplication load balancing scenario, you must upgrade the media servers to NetBackup 8.1.1 or later and modify the `pd.conf` file on all the media servers. If a backup job selects an older media server (earlier than NetBackup 8.1.1) for the load balancing pool, fixed-length deduplication is used instead of variable-length deduplication. Avoid configuring media servers with different NetBackup versions in a load balancing scenario. The data segments generated from variable-length deduplication are different from the data segments generated from fixed-length deduplication. Therefore, load balancing media servers with different NetBackup versions results in a low deduplication ratio.

See [“About the MSDP `pd.conf` configuration file”](#) on page 209.

See [“Editing the MSDP `pd.conf` file”](#) on page 210.

See [“MSDP `pd.conf` file parameters”](#) on page 210.

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/openv/lib/ost-plugins/`

- (Windows) `install_path\Veritas\NetBackup\bin\ost-plugins`

See “MSDP pd.conf file parameters” on page 210.

See “Editing the MSDP pd.conf file” on page 210.

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: Veritas recommends that you make a backup copy of the file before you edit it.

See “About the MSDP pd.conf configuration file” on page 209.

See “MSDP pd.conf file parameters” on page 210.

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 5-36](#) 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:

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

Table 5-36 `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 during backups.</p> <p>By default, the data is compressed.</p> <p>Default value: <code>COMPRESSION = 1</code></p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About MSDP compression” on page 127.</p>

Table 5-36 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, Veritas recommends that you do not enable this option.■ On high load systems, the load may delay the capacity information reporting. If so, NetBackup may mark the storage unit as down. <p>Default value: <code>CR_STATS_TIMER = 0</code></p> <p>Possible values: 0 or greater, in seconds</p>
DEBUGLOG	<p>Specifies the file to which NetBackup writes the deduplication 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: <code>DEBUGLOG = /var/log/puredisk/pdplugin.log</code>■ Windows: <code>DEBUGLOG = C:\pdplugin.log</code> <p>Possible values: Any path</p>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
DISABLE_BACKLEVEL_TLS	<p>When secure communication is established between the client and the server, this parameter specifies whether or not to disable older TLS versions. NetBackup version 8.0 and earlier use older TLS versions such as SSLV2, SSLV3, TLS 1.0, and TLS 1.1.</p> <p>To enable TLS 1.2, change the value of the DISABLE_BACKLEVEL_TLS parameter to 1 and restart the NetBackup Deduplication Engine (spoold) and the NetBackup Deduplication Manager (spad).</p> <p>Default value: <code>DISABLE_BACKLEVEL_TLS = 0</code></p> <p>Possible values: 0 (off) or 1 (on)</p> <p>Note: To enable TLS 1.2, NetBackup version must be 8.1 and later. When TLS 1.2 is enabled (<code>DISABLE_BACKLEVEL_TLS = 1</code>) on a machine (which can be a client or a media server or a load balance server), to establish communication, all machines connected to it must also enable TLS 1.2.</p> <p>For a standard backup, NetBackup client version 8.0 and earlier can communicate with NetBackup server (media server or load balance server) version 8.1 that has TLS 1.2 enabled.</p> <p>However, in case of optimized duplication and replication, load balance, and client direct duplication, NetBackup client versions 8.0 and earlier cannot communicate with NetBackup server (media server or load balance server) version 8.1, which has TLS 1.2 enabled.</p> <p>Further, NetBackup 5000 series appliance does not support TLS 1.2.</p>
DONT_SEGMENT_TYPES	<p>A comma-separated list of file name extensions of files not to be deduplicated. Files in the backup stream that have the specified extensions are given a single segment if smaller than 16 MB. Larger files are deduplicated using the maximum 16-MB segment size.</p> <p>Example: <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>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
ENCRYPTION	<p>Specifies whether to encrypt the data during backups. By default, files are not encrypted.</p> <p>If you set this parameter to 1 on all hosts, the data is encrypted during transfer and on the storage.</p> <p>Default value: <code>ENCRYPTION = 0</code></p> <p>Possible values: 0 (no encryption) or 1 (encryption)</p> <p>See "About MSDP encryption" on page 129.</p>
FIBRECHANNEL	<p>Enable Fibre Channel for backup and restore traffic to and from a NetBackup series appliance.</p> <p>Default value: <code>FIBRECHANNEL = 0</code></p> <p>Possible values: 0 (off) or 1 (on)</p>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
FILE_KEEP_ALIVE_INTERVAL	<p>The interval in seconds at which to perform keepalives on idle sockets.</p> <p>The following items describe the behavior based on how you configure this parameter:</p> <ul style="list-style-type: none">■ 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 200.■ 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> <p>Using keepalive interval of xxxx seconds</p> <p>For more information about the deduplication plug-in log file, see <code>DEBUGLOG</code> and <code>LOGLEVEL</code> in this table.</p>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
FP_CACHE_CLIENT_POLICY	<p>Note: Veritas 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 “Configuring MSDP fingerprint cache seeding on the client” on page 72.</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 Veritas 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 5-36 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: FP_CACHE_MAX_COUNT = 1024</p> <p>Possible values: 0 to 4096</p> <p>Note: Change this value only when directed to do so by a Veritas representative.</p>
FP_CACHE_MAX_MBSIZE	<p>Specifies the amount of memory in MBs to use for the fingerprint cache.</p> <p>Default value: FP_CACHE_MAX_MBSIZE = 20</p> <p>Possible values: 0 to the computer limit</p> <p>Note: Change this value only when directed to do so by a Veritas 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 FP_CACHE_PERIOD_REBASING_THRESHOLD value. <p>Default value: FP_CACHE_PERIOD_REBASING_THRESHOLD = 16</p> <p>Possible values: 0 (disabled) to 256</p> <p>See “About MSDP storage rebasing” on page 368.</p>
FP_CACHE_REBASING_THRESHOLD	<p>Specifies the threshold (MB) for normal rebasing during backups. A container is considered for rebasing if both of the following are true:</p> <ul style="list-style-type: none"> ■ 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: FP_CACHE_REBASING_THRESHOLD = 4</p> <p>Possible values: 0 (disabled) to 200</p> <p>If you change this value, consider the new value carefully. If you set it too large, all containers become eligible for rebasing. Deduplication rates are lower for the backup jobs that perform rebasing.</p> <p>See “About MSDP storage rebasing” on page 368.</p>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
LOCAL_SETTINGS	<p>Specifies whether to use the <code>pd.conf</code> settings of the local host or to allow the server to override the local settings. The following is the order of precedence for local settings:</p> <ul style="list-style-type: none"> ■ Local host ■ Load balancing server ■ Storage server <p>To use the local settings, set this value to 1.</p> <p>Default value: <code>LOCAL_SETTINGS = 0</code></p> <p>Possible values: 0 (allow override) or 1 (always use local settings)</p>
LOGLEVEL	<p>Specifies the amount of information that is written to the log file. The range is from 0 to 10, with 10 being the most logging.</p> <p>Default value: <code>LOGLEVEL = 0</code></p> <p>Possible values: An integer, 0 to 10 inclusive</p> <p>Note: Change this value only when directed to do so by a Veritas representative.</p>
MAX_IMG_MBSIZE	<p>The maximum backup image fragment size in megabytes.</p> <p>Default value: <code>MAX_IMG_MBSIZE = 51200</code></p> <p>Possible values: 0 to 51,200, in MBs</p> <p>Note: Change this value only when directed to do so by a Veritas representative.</p>
MAX_LOG_MBSIZE	<p>The maximum size of the log file in megabytes. NetBackup creates a new log file when the log file reaches this limit. NetBackup prepends the date and the ordinal number beginning with 0 to each log file, such as <code>120131_0_pdplugin.log</code>, <code>120131_1_pdplugin.log</code>, and so on.</p> <p>Default value: <code>MAX_LOG_MBSIZE = 100</code></p> <p>Possible values: 0 to 50,000, in MBs</p>
META_SEGKSIZE	<p>The segment size for metadata streams</p> <p>Default value: <code>META_SEGKSIZE = 16384</code></p> <p>Possible values: 32-16384, multiples of 32</p> <p>Note: Change this value only when directed to do so by a Veritas representative.</p>

Table 5-36 `pd.conf` file parameters (*continued*)

Parameter	Description
MTSTRM_BACKUP_CLIENTS	<p>If set, limits the use of the Multi-Threaded Agent to the backups of the specified clients. The clients that are not specified use single-threading.</p> <p>This setting does not guarantee that the specified clients use the Multi-Threaded Agent. The <code>MaxConcurrentSessions</code> parameter in the <code>mtstrm.conf</code> file controls the number of backups the Multi-Threaded Agent processes concurrently. If you specify more clients than the <code>MaxConcurrentSessions</code> value, some of the clients may use single-threaded processing.</p> <p>See “MSDP <code>mtstrm.conf</code> file parameters” on page 62.</p> <p>The format is a comma-separated list of the clients, case insensitive (for example, <code>MTSTRM_BACKUP_CLIENTS = client1,client2,client3</code>).</p> <p>Default value: <code>MTSTRM_BACKUP_CLIENTS = (no default value)</code></p> <p>Possible values: comma separated client names</p> <p>See “About the MSDP Deduplication Multi-Threaded Agent” on page 59.</p>

Table 5-36 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: MTSTRM_BACKUP_ENABLED = (no default value)</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 MTSTRM_BACKUP_ENABLED 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 (MTSTRM_BACKUP_ENABLED = 1).■ 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 (MTSTRM_BACKUP_ENABLED = 0). <p>See "About the MSDP Deduplication Multi-Threaded Agent" on page 59.</p>

Table 5-36 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 59.</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 59.</p>
OPTDUP_BANDWIDTH	<p>Determines the bandwidth that is allowed for each optimized duplication and Auto Image Replication stream on a deduplication server. <code>OPTDUP_BANDWIDTH</code> does not apply to clients. The value is specified in KBytes/second.</p> <p>Default value: <code>OPTDUP_BANDWIDTH= 0</code></p> <p>Possible values: 0 (no limit) to the practical system limit, in KBs/sec</p> <p>A global bandwidth parameter effects whether or not <code>OPTDUP_BANDWIDTH</code> applies.</p> <p>See “About configuring MSDP optimized duplication and replication bandwidth” on page 191.</p>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
OPTDUP_COMPRESSION	<p>Specifies whether to compress the data during optimized duplication and Auto Image Replication. By default, files are compressed. To disable compression, change the value to 0. This parameter does not apply to clients.</p> <p>Default value: OPTDUP_COMPRESSION = 1</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About MSDP compression” on page 127.</p>
OPTDUP_ENCRYPTION	<p>Specifies whether to encrypt the data during optimized duplication and replication. By default, files are not encrypted. If you want encryption, change the value to 1 on the MSDP storage server and on the MSDP load balancing servers. This parameter does not apply to clients.</p> <p>If you set this parameter to 1 on all hosts, the data is encrypted during transfer.</p> <p>Default value: OPTDUP_ENCRYPTION = 0</p> <p>Possible values: 0 (off) or 1 (on)</p> <p>See “About MSDP encryption” on page 129.</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>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
PDDO_SSL_VERIFY	<p>Specifies whether to perform SSL verification when communicating with the NetBackup 5000 series appliance. By default, SSL verification is not performed while communication with the NetBackup 5000 series appliance.</p> <p>To perform SSL verification, change the value to 1 and copy the root certificate file of NetBackup 5000 series appliance to following directory:</p> <ul style="list-style-type: none">■ For Windows: install_path\Veritas\NetBackup\bin\ost-plugins■ For UNIX: /usr/opensv/lib/ost-plugins/ <p>After copying the root certificate file, you must rename it to storage_server_name.pem, where <i>storage_server_name</i> is the fully qualified domain name of the NetBackup 5000 series appliance.</p> <p>The root certificate file on NetBackup 5000 series appliance is:</p> <ul style="list-style-type: none">■ (1.4.4) /Storage/var/VRTSat/.VRTSat/profile/certstore/trusted/*.0■ (1.4.5) /Storage/var/VxAT/root/.VRTSat/profile/certstore/trusted/*.0 <p>If there are multiple NetBackup servers and NetBackup clients that communicate with NetBackup 5000 series appliance, the root certificate file must be copied to all the media servers and NetBackup clients.</p> <p>Default value: PDDO_SSL_VERIFY = 0</p> <p>Possible values: 0 (off) or 1 (on)</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: edb are Exchange Server files; md fare SQL Server master database files, ndf are SQL Server secondary data files, and segsize64k are Microsoft SQL streams.</p> <p>Default value: PREFERRED_EXT_SEGKSIZE = edb:32,mdf:64,ndf:64,segsize64k:64</p> <p>Possible values: <i>file_extension:segment_size_in_KBs</i> pairs, separated by commas.</p> <p>See also SEGKSIZE.</p>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
PREFETCH_SIZE	<p>The size in bytes to use for the data buffer for restore operations.</p> <p>Default value: 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 Veritas representative.</p>
RESTORE_DECRYPT_LOCAL	<p>Specifies on which host to decrypt and decompress the data during restore operations.</p> <p>Depending on your environment, decryption and decompression on the client may provide better performance.</p> <p>Default value: RESTORE_DECRYPT_LOCAL = 1</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 Veritas representative.</p> <p>You can also specify the segment size for specific file types. See PREFERRED_EXT_SEGKSIZE.</p>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
VLD_CLIENT_NAME	<p>Specifies the name of the NetBackup client to enable variable-length deduplication. By default, the <code>VLD_CLIENT_NAME</code> parameter is not present in the <code>pd.conf</code> configuration file.</p> <p>You can also specify different maximum and minimum segment sizes with this parameter for different NetBackup clients. If you do not specify the segment sizes, then the default values are considered.</p> <p>The values are case-sensitive.</p> <p>Use in any of the following formats:</p> <ul style="list-style-type: none">■ <code>VLD_CLIENT_NAME = *</code> Enables variable-length deduplication for all NetBackup clients and uses the default <code>VLD_MIN_SEGKSIZE</code> and <code>VLD_MAX_SEGKSIZE</code> values.■ <code>VLD_CLIENT_NAME = clientname</code> Enables variable-length deduplication for NetBackup client <code>clientname</code> and uses the default <code>VLD_MIN_SEGKSIZE</code> and <code>VLD_MAX_SEGKSIZE</code> values.■ <code>VLD_CLIENT_NAME = clientname (64, 256)</code> Enables variable-length deduplication for NetBackup client <code>clientname</code> and uses 64 KB as the <code>VLD_MIN_SEGKSIZE</code> and 256 KB as the <code>VLD_MAX_SEGKSIZE</code> value. <p>Note: You can add a maximum of 50 clients in the <code>pd.conf</code> file.</p>
VLD_MIN_SEGKSIZE	<p>The minimum size of the data segment for variable-length deduplication in KB. The segment size must be in multiples of 4 and fall in between 4 KB to 16384 KB. The default value is 64 KB.</p> <p>The value must be smaller than <code>VLD_MAX_SEGKSIZE</code>. Different NetBackup clients can have different segment sizes.</p> <p>A larger value reduces the CPU consumption, but decreases the deduplication ratio. A smaller value increases the CPU consumption, but increases the deduplication ratio</p> <p>Note: Keeping similar or close values for <code>VLD_MIN_SEGKSIZE</code> and <code>VLD_MAX_SEGKSIZE</code> results in a performance that is similar to fixed-length deduplication.</p>

Table 5-36 pd.conf file parameters (*continued*)

Parameter	Description
VLD_MAX_SEGKSIZE	<p>The maximum size of the data segment for variable-length deduplication in KB. VLD_MAX_SEGKSIZE is used to set a boundary for the data segments. The segment size must be in multiples of 4 and fall in between 4 KB to 16384 KB. The default value is 128 KB.</p> <p>The value must be greater than VLD_MIN_SEGKSIZE. Different NetBackup clients can have different segment sizes.</p> <p>Note: Keeping similar or close values for VLD_MIN_SEGKSIZE and VLD_MAX_SEGKSIZE results in a performance that is similar to fixed-length deduplication.</p>
VLD_POLICY_NAME	<p>Specifies the name of the backup policy to enable variable-length deduplication. By default, the VLD_POLICY_NAME parameter is not present in the pd.conf configuration file.</p> <p>You can also specify different maximum and minimum segment sizes with this parameter for different NetBackup policies. If you do not specify the segment sizes, then the default values are considered.</p> <p>The values are case-sensitive.</p> <p>Use in any of the following formats:</p> <ul style="list-style-type: none">■ VLD_POLICY_NAME = *■ VLD_POLICY_NAME = <i>polycyname</i>■ VLD_POLICY_NAME = <i>polycyname</i> (64, 256) <p>Note: You can add a maximum of 50 policies in the pd.conf file.</p>

See [“About the MSDP pd.conf configuration file”](#) on page 209.

See [“Editing the MSDP pd.conf file”](#) on page 210.

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 Veritas 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 Veritas 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, Veritas recommends that you get the storage server configuration and save it in a file. The file does not exist unless you create it.

The following is an example of a populated configuration file:

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

```
V7.0 "Catalog files Count" "1288" string
V7.0 "Space Used Within Containers" "142.3GB" string
```

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

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

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

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

See [“Saving the MSDP storage server configuration”](#) on page 228.

See [“Editing an MSDP storage server configuration file”](#) on page 229.

See [“Setting the MSDP storage server configuration”](#) on page 230.

Saving the MSDP storage server configuration

Veritas 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 227.

See [“Recovering from an MSDP storage server disk failure”](#) on page 380.

See [“Recovering from an MSDP storage server failure”](#) on page 381.

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 5-37](#) shows the configuration lines that are required.

Table 5-37 Required lines for a recovery file

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

Table 5-37 Required lines for a recovery file (continued)

Configuration setting	Description
V7.0 "encryption" " " int	The value should be the same as the value that was used when you configured the storage server.
V7.0 "kmsenabled" " " int	The value is used to enable or disable MSDP KMS configuration. The value should be the same as the value that was used when you configured the storage server.
V7.0 "kmsservertype" " " int	The value is KMS server type. This value should be 0.
V7.0 "kmsservername" " " string	<p>The value is NBU Key Management Server. The value should be the same as the value that was used when you configured the storage server.</p> <p>If you use an external KMS as a KMS server, the value must be the NetBackup master server name. See <i>External KMS support in NetBackup</i> in the <i>NetBackup Security and Encryption Guide</i>.</p>
V7.0 "keygroupname" " " string	The value should be the same as the value that was used when you configured the storage server.

See [“About saving the MSDP storage server configuration”](#) on page 227.

See [“Recovering from an MSDP storage server disk failure”](#) on page 380.

See [“Recovering from an MSDP storage server failure”](#) on page 381.

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 228.

- 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 5-37](#)) 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 380.

See [“Recovering from an MSDP storage server failure”](#) on page 381.

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

See [“About saving the MSDP storage server configuration”](#) on page 227.

See [“Saving the MSDP storage server configuration”](#) on page 228.

See [“Editing an MSDP storage server configuration file”](#) on page 229.

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.example.com`, the configuration file name is `DedupeServer.example.com.cfg`.

The following is the location of the file:

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

UNIX: `/usr/opensv/lib/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 231.

To delete the host configuration file

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

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

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

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

`DedupeServer.example.com.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 343.

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/opensv/pdde/pdconfigure/cfg/userconfigs/pdregistry.cfg
/etc/pdregistry.cfg
```

To reset the MSDP registry on Windows

- ◆ Delete the contents of the following keys in the Windows registry:
 - `HKLM\SOFTWARE\Symantec\PureDisk\Agent\ConfigFilePath`

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

Warning: Editing the Windows registry may cause unforeseen results.

About protecting the MSDP catalog

To increase availability, NetBackup provides a two-tier approach to protect the MSDP catalog, as follows:

- | | |
|-----------------------|--|
| Daily shadow copies | NetBackup automatically creates copies of the MSDP catalog.
See “About the MSDP shadow catalog” on page 233. |
| Catalog backup policy | Veritas provides a utility that you can use to configure a NetBackup policy that backs up the MSDP catalog.
See “About the MSDP catalog backup policy” on page 234. |

See [“About recovering the MSDP catalog”](#) on page 377.

About the MSDP shadow catalog

The NetBackup Deduplication Manager automatically creates a *shadow copy* of the catalog daily. The Deduplication Manager also builds a transaction log for each shadow copy. If NetBackup detects corruption in the MSDP catalog, the Deduplication Manager restores the catalog automatically from the most recent shadow copy. That restore process also plays the transaction log so that the recovered MSDP catalog is current.

By default, the NetBackup Deduplication Manager stores the shadow copies on the same volume as the catalog itself. Veritas recommends that you store the shadow copies on a different volume.

Warning: You can change the path only during initial MSDP configuration only. If you change it after MSDP backups exist, data loss may occur.

See [“Changing the MSDP shadow catalog path”](#) on page 236.

The NetBackup Deduplication Manager creates a shadow copy at 0340 hours daily, host time. To change the schedule, you must change the scheduler definition file.

See [“Changing the MSDP shadow catalog schedule”](#) on page 237.

By default, the NetBackup Deduplication Manager keeps five shadow copies of the catalog. You can change the number of copies.

See [“Changing the number of MSDP catalog shadow copies”](#) on page 238.

About the MSDP catalog backup policy

Veritas recommends that you protect the MSDP catalog by backing it up. (A NetBackup catalog backup does not include the MSDP catalog.) The NetBackup Deduplication Catalog Policy Administration and Catalog Disaster Recovery utility (the `drcontrol` utility) configures a backup policy for the MSDP catalog. The policy also includes other important MSDP configuration information.

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

The following are the attributes for the catalog backup policy that the `drcontrol` utility creates:

Schedule	Weekly Full Backup and daily Differential Incremental Backup .
Backup window	6:00 A.M. to 6:00 P.M.
Retention	2 weeks

Backup selection The following are the default catalog paths.

UNIX:

```
/database_path/databases/catalogshadow
/storage_path/etc
/database_path/databases/spa
/storage_path/var
/usr/opensv/lib/ost-plugins/pd.conf
/usr/opensv/lib/ost-plugins/mtstrm.conf
/database_path/databases/datacheck
```

Additional catalog paths are backed up if using a Cloud Catalyst storage server:

```
/msdpc/storage/databases/catalog/___info___2
/msdpc/storage/databases/catalog/2
/msdpc/storage/data/.dataconverted
/msdpc/storage/data/.format
/msdpc/storage/data/.gcnf
/msdpc/storage/data/.identity
/msdpc/storage/data/.srcmap
/msdpc/cache/etc
/msdpc/cache/fsdb
/msdpc/cache/userdata/proc/cloud.lsu
```

Windows:

```
database_path\databases\catalogshadow
storage_path\etc
storage_path\var
install_path\Veritas\NetBackup\bin\ost-plugins\pd.conf
install_path\Veritas\NetBackup\bin\ost-plugins\mtstrm
database_path\databases\spa
database_path\databases\datacheck
```

By default, NetBackup uses the same path for the storage and the catalog; the `database_path` and the `storage_path` are the same. If you configure a separate path for the deduplication database, the paths are different. Regardless, the `drcontrol` utility captures the correct paths for the catalog backup selections.

You should consider the following items carefully before you configure an MSDP catalog backup:

- Do not use the **Media Server Deduplication Pool** as the destination for the catalog backups. Recovery of the MSDP catalog from its **Media Server Deduplication Pool** is impossible.

- Use a storage unit that is attached to a NetBackup host other than the MSDP storage server.
- Use a separate MSDP catalog backup policy for each MSDP storage server. The `drcontrol` utility does not verify that the backup selections are the same for multiple storage servers. If the backup policy includes more than one MSDP storage server, the backup selection is the union of the backup selections for each host.
- You cannot use one policy to protect MSDP storage servers on both UNIX hosts and Windows hosts.
UNIX MSDP storage servers require a Standard backup policy and Windows MSDP storage servers require an MS-Windows policy.

See [“Configuring an MSDP catalog backup”](#) on page 239.

See [“Updating an MSDP catalog backup policy”](#) on page 244.

Changing the MSDP shadow catalog path

You can change the location of the catalog shadow copies. It is recommended that you store the copies on a different volume than both the `storage_path` and the `database_path`. (If you configured a separate path for the deduplication database, the paths are different.)

NetBackup stores the MSDP catalog shadow copies in the following location:

UNIX: `/database_path/databases/catalogshadow`

Windows: `database_path\databases\catalogshadow`

Warning: You can change the shadow catalog path during initial MSDP configuration only. If you change it after MSDP backups exist, data loss may occur.

See [“About protecting the MSDP catalog”](#) on page 233.

To change the MSDP catalog shadow path

- 1 Open the following file in a text editor:
UNIX: `/storage_path/etc/puredisk/spa.cfg`
Windows: `storage_path\etc\puredisk\spa.cfg`
- 2 Find the `CatalogShadowPath` parameter and change the value to the wanted path.
The volume must be mounted and available.

- 3 After your changes, save the file.
- 4 Restart the NetBackup Deduplication Manager (*spad*).
- 5 Create the shadow catalog directories by invoking the following command on the MSDP storage server:

UNIX: `/usr/opensv/pdde/pdcr/bin/cacontrol --catalog backup all`

Windows: `install_path\Veritas\pdde\cacontrol --catalog backup all`

- 6 If an MSDP catalog backup policy exists, update the policy with the new shadow catalog directories. To do so, invoke the following command on the MSDP storage server:

UNIX: `/usr/opensv/pdde/pdcr/bin/drcontrol --update_policy --policy policy_name`

Windows: `install_path\Veritas\pdde\drcontrol --update_policy --policy policy_name`

Changing the MSDP shadow catalog schedule

NetBackup automatically creates a copy of the MSDP catalog at 0340 hours daily, host time. You can change the default schedule.

See [“About protecting the MSDP catalog”](#) on page 233.

To change the MSDP shadow catalog schedule

- 1** Open the following file in a text editor:

UNIX: `/database_path/databases/spa/database/scheduler/5`

Windows: `database_path\databases\spa\database\scheduler\5`

By default, NetBackup uses the same path for the storage and the catalog; the `database_path` and the `storage_path` are the same. If you configure a separate path for the deduplication database, the paths are different.

The contents of the file are similar to the following line. The second section of the line (40 3 * * *) configures the schedule.

```
CatalogBackup|40 3 * * *|21600|32400|
```

- 2** Edit the second section of the file (`40 3 * * *`). The schedule section conforms to the UNIX `crontab` file convention, as follows:

40 3 * * *

T T T T T

```
| | | | _____ Day of week (0 - 7, Sunday is both 0 and 7, or use  
| | | | sun, mon, tue, wed, thu, fri, sat; asterisk (*) is  
| | | | every day)  
| | | | _____ Month (1 - 12; asterisk (*) is every month)  
| | | | _____ Day of month (1 - 31; asterisk (*) is every  
| | | | day of the month)  
| | | | _____ Hour (0 - 23; asterisk (*) is every hour)  
| | | | _____ Minute (0 - 59; asterisk (*) is every  
| | | | minute of the hour)
```

- 3 After your changes, save the file.
- 4 Restart the NetBackup Deduplication Manager (`spad`).

Changing the number of MSDP catalog shadow copies

NetBackup keeps five shadow copies of the MSDP catalog. You can change the number of copies.

See “About protecting the MSDP catalog” on page 233.

To change the number of MSDP catalog shadow copies

- 1 Open the following file in a text editor:
UNIX: `/storage_path/etc/puredisk/spa.cfg`
Windows: `storage_path\etc\puredisk\spa.cfg`
- 2 Find the `CatalogBackupVersions` parameter and change the value to the wanted number of shadow copies. The valid values are 1 to 256, inclusive.
- 3 After your changes, save the file.
- 4 Restart the NetBackup Deduplication Manager (`spad`).

Configuring an MSDP catalog backup

Use the following procedure to configure a backup policy for the NetBackup MSDP catalog.

See [“About protecting the MSDP data”](#) on page 45.

See [“Troubleshooting MSDP catalog backup”](#) on page 429.

To configure an MSDP catalog backup

- 1 Verify that the MSDP storage server host (that is, the media server) is an additional server for the NetBackup master server. See **NetBackup Management > Host Properties > *masterserver_name* > Servers > Additional Servers** in the NetBackup Administration Console.

If the storage server is not in the **Additional Servers** list, add the MSDP storage server host to the **Additional Servers** list. The host *must* be in the **Additional Servers** list and *cannot* be in the **Media Servers** list.

- 2 On the MSDP storage server, invoke the `drcontrol` utility and use the appropriate options for your needs. The following is the syntax for the utility:

Windows: `install_path\Veritas\pdde\drcontrol --new_policy --residence residence [--policy policy_name] [--client host_name] [--hardware machine_type] [--OS operating_system] [--dsid data_selection_ID] [--NB_install_dir install_directory]`

UNIX: `/usr/opensv/pdde/pdcr/bin/drcontrol --new_policy --residence residence [--policy policy_name] [--disk_pool disk_pool_name] [--client host_name] [--hardware machine_type] [--OS operating_system] [--dsid data_selection_ID]`

Descriptions of the options are available in another topic. Note: To ensure that NetBackup activates the policy, you must specify the `--residence residence` option.

See “MSDP `drcontrol` options” on page 240.

The utility creates a log file and displays its path in the command output.

See “NetBackup MSDP log files” on page 407.

MSDP `drcontrol` options

The `drcontrol` utility resides in the following directories, depending on host type:

- UNIX: `/usr/opensv/pdde/pdcr/bin`
- Windows: `install_path\Veritas\pdde`

The `drcontrol` utility creates a log file.

See “NetBackup MSDP log files” on page 407.

[Table 5-38](#) describes the options for creating and updating an MSDP catalog backup policy.

Table 5-38 MSDP `drcontrol` options for catalog backup and recovery

Option	Description
<code>--auto_recover_DR</code>	<p>Recover the MSDP catalog from the most recent backup image. This option automatically recovers the catalog and performs all of the actions necessary to return MSDP to full functionality.</p> <p>This option requires the <code>--policy <i>policy_name</i></code> option.</p> <p>To recover the catalog from a backup other than the most recent, contact your Veritas Support representative.</p>
<code>--client <i>host_name</i></code>	<p>The client to back up (that is, the host name of the MSDP storage server).</p> <p>Default: the value that <code>bpgetconfig CLIENT_NAME</code> returns.</p>
<code>--cleanup</code>	<p>Remove all of the old MSDP catalog directories during the catalog recovery process. Those directories are renamed during the recovery.</p>
<code>--delete_old_files_for_cloud_catalyst</code>	<p>Deletes the existing deduplication catalog files before restoring them from the last backup in the <code>recover_last_image</code> step. Used for a cloud media server after the file system has been recovered in the <code>recover_last_cloud_catalyst_image</code> step.</p>
<code>--disk_pool</code>	<p>In a NetBackup Cloud Catalyst recovery scenario, you must specify the disk pool name that is configured for the cloud storage server. This option is required for <code>auto_recover_DR</code> when the disk pool name cannot be determined from the host name. For example, for a cloud storage server that is configured for Cloud Catalyst.</p>
<code>--dsid</code>	<p>The data selection ID is the catalog directory for one of the NetBackup domains.</p> <p>In a multi-domain scenario when you recover the catalog from another domain, the <code>dsid</code> of the other NetBackup domain is used. To obtain the <code>dsid</code> of the other NetBackup domain, run the <code>spauser</code> command to list the <code>dsid</code>.</p> <p>The default value is 2.</p>
<code>--hardware <i>machine_type</i></code>	<p>The hardware type or the computer type for the host.</p> <p>Spaces are not allowed. If the string contains special characters, enclose it in double quotation marks ("").</p> <p>Default: Unknown.</p>

Table 5-38 MSDP `drcontrol` options for catalog backup and recovery
(continued)

Option	Description
<code>--initialize_DR</code>	Performs the following actions to prepare for MSDP catalog recovery: <ul style="list-style-type: none"> ■ Verifies that the most recent catalog backup is valid. ■ Stops the deduplication services. ■ Moves the existing catalog files so that they are empty for the recovery.
<code>--list_files</code>	List the files in the most recent MSDP catalog backup. This option requires the <code>--policy policy_name</code> option.
<code>--log_file pathname</code>	The pathname for the log file that the <code>drcontrol</code> utility creates. By default, the utility writes log files to <code>/storage_path/log/drcontrol/</code> .
<code>--NB_install_dir install_directory</code>	Windows only. Required option if NetBackup was installed in a location other than the default (C:\Program Files\Veritas). If the string contains spaces or special characters, enclose it in double quotation marks ("). Do not use a trailing backslash in the <code>install_directory</code> string.
<code>--new_policy</code>	Create a new policy to protect the deduplication catalog on this host. If a policy with the given name exists already, the command fails. Note: To ensure that NetBackup activates the policy, you must specify the <code>--residence residence</code> option.
<code>--OS operating_system</code>	The operating system for the host. Spaces are not allowed. If the string contains special characters, enclose it in double quotation marks ("). Default: UNIX/Linux or MS-Windows.
<code>--policy policy_name</code>	The name for the backup policy. Required with <code>--auto_recover_DR</code> and <code>--update_policy</code> ; optional with <code>--new_policy</code> . Default: <code>Dedupe_Catalog_shorthostname</code>
<code>--print_space_required</code>	Display an estimate of the percentage of file system space that is required to recover the MSDP catalog.

Table 5-38 MSDP `drcontrol` options for catalog backup and recovery
(continued)

Option	Description
<code>--recover_last_cloud_catalyst_image</code>	Use <code>bprestore</code> to restore the cloud media server-related contents of the last deduplication catalog backup. This operation requires that disaster recovery be initialized.
<code>--recover_last_image</code>	Restore the MSDP catalog from the last set of backup images (that is, the last full plus all subsequent incrementals). The <code>drcontrol</code> utility calls the NetBackup <code>bprestore</code> command for the restore operation.
<code>--refresh_shadow_catalog</code>	Deletes all existing shadow catalog copies and creates a new catalog shadow copy.
<code>--residence <i>residence</i></code>	<p>The name of the storage unit on which to store the MSDP catalog backups.</p> <p>Do not use the Media Server Deduplication Pool as the destination for the catalog backups. Recovery of the MSDP catalog from its Media Server Deduplication Pool is impossible.</p> <p>Veritas recommends that you use a storage unit that is attached to a NetBackup host other than the MSDP storage server.</p>

Table 5-38 MSDP `drcontrol` options for catalog backup and recovery
(continued)

Option	Description
<code>--update_policy</code>	<p>Update a policy, as follows:</p> <ul style="list-style-type: none">■ If the client name (of this media server) is not in the policy's client list, add the client name to the policy's client list.■ If you specify the <code>--os</code> or <code>--hardware</code> options, replace the values currently in the policy with the new values.■ Update the backup selection based on the locations of the MSDP storage directories and configuration files. Therefore, if you modify any of the following, you must use this option to update the catalog backup policy:<ul style="list-style-type: none">■ Any of the following values in the <code>spa.cfg</code> file (section:variable pairs):<ul style="list-style-type: none">■ <code>StorageDatabase:CatalogShadowPath</code>■ <code>StorageDatabase:Path</code>■ <code>Paths:Var</code>■ The <code>spa.cfg</code> or <code>contentrouter.cfg</code> locations in the <code>pdregistry.cfg</code> file. <p>This option fails if there is no policy with the given policy name. It also fails if the existing policy type is incompatible with the operating system of the host on which you run the command.</p> <p>This option requires the <code>--policy policy_name</code> option.</p>
<code>--verbose</code>	Echo all <code>drcontrol</code> log statements to stdout.

See [“Configuring an MSDP catalog backup”](#) on page 239.

Updating an MSDP catalog backup policy

You can use any NetBackup method to update an MSDP catalog backup policy manually. However, you should use the NetBackup Deduplication Catalog Policy Administration and Catalog Disaster Recovery (`drcontrol`) under the following circumstances:

- To add the client name of the storage server to the policy's client list.
- To update the `--os` value.
- To update the `--hardware` value.
- To update the backup selection if you modified any of the following configuration values:

- Any of the following values in the `spa.cfg` file (section:variable pairs):
 - `StorageDatabase:CatalogShadowPath`
 - `StorageDatabase:Path`
 - `Paths:Var`
- The `spa.cfg` or `contentrouter.cfg` locations in the `pdregistry.cfg` file.

See “About protecting the MSDP data” on page 45.

See “Troubleshooting MSDP catalog backup” on page 429.

To update an MSDP catalog backup

- ◆ On the MSDP storage server, invoke the `drcontrol` utility and use the appropriate options for your needs. The following is the syntax for an update operation:

UNIX: `/usr/opensv/pdde/pdcr/bin/drcontrol --update_policy --policy policy_name [--client host_name] [--hardware machine_type] [--OS operating_system]`

Windows: `install_path\Veritas\pdde\drcontrol --update_policy --policy policy_name [--client host_name] [--hardware machine_type] [--OS operating_system] [--OS operating_system] [--NB_install_dir install_directory]`

Descriptions of the options are available in another topic.

See “MSDP `drcontrol` options” on page 240.

The utility creates a log file and displays its path in the command output.

See “NetBackup MSDP log files” on page 407.

About MSDP FIPS compliance

The Federal Information Processing Standards (FIPS) define U.S. and Canadian Government security and interoperability requirements for computer systems. The FIPS 140-2 standard specifies the security requirements for cryptographic modules. It describes the approved security functions for symmetric and asymmetric key encryption, message authentication, and hashing.

For more information about the FIPS 140-2 standard and its validation program, see the National Institute of Standards and Technology (NIST) and the Communications Security Establishment Canada (CSEC) Cryptographic Module Validation Program website at

<https://csrc.nist.gov/projects/cryptographic-module-validation-program>.

The NetBackup MSDP is now FIPS validated and can be operated in FIPS mode.

Note: You must run FIPS mode on a new installation of NetBackup 8.1.1.

Enabling MSDP FIPS mode

Ensure that you configure the storage server before you enable the MSDP FIPS mode.

Caution: Enabling MSDP FIPS mode might affect the NetBackup performance on a server with the Solaris operating system.

Enable the FIPS mode for MSDP by running the following commands:

- For UNIX:

```
/usr/opensv/pdde/pdag/scripts/set_fips_mode.sh 1
```
- For Windows:

```
<install_path>\Veritas\pdde\set_fips_mode.bat 1
```
- Restart the NetBackup service on the server and the client.
 - For UNIX:
 - ```
/usr/opensv/netbackup/bin/bp.kill_all
```
    - ```
/usr/opensv/netbackup/bin/bp.start_all
```
 - For Windows:
 - ```
<install_path>\NetBackup\bin\bpdown
```
    - ```
<install_path>\NetBackup\bin\bpup
```

Warning: It is recommended that you do not disable the MSDP FIPS mode once you enable it, for security reasons.

Getting the status of MSDP FIPS mode

To get status of the MSDP FIPS mode, enter the following commands:

For UNIX:

```
/usr/opensv/pdde/pdcr/bin/crcontrol --getmode
```

For Windows:

```
<install_path>\Veritas\pdde\crcontrol.exe --getmode
```

Other things to note:

- FIPS must be enabled on all the NetBackup components to establish a connection. When the FIPS mode is not enabled, communication can occur between the NetBackup clients and the servers that have earlier, supported NetBackup versions.
- The MSDP FIPS mode does not support NetBackup Cloud Catalyst.

Configuring the NetBackup client-side deduplication to support multiple interfaces of MSDP

If you have network configurations like VLAN or Subnet with NetBackup client, then the MSDP server has multiple network interfaces. These interfaces are connected to different switches or VLANs. As MSDP has only one storage server, NetBackup clients cannot access the MSDP server using the storage server name. The deduplication can fail on the clients.

You can add support for up to 30 interfaces.

Run the following steps to use the `cacontrol` command option (Location: `/usr/openv/pdde/pdcr/bin/`) to configure MSDP and specify the network interfaces that the NetBackup client can use:

- 1 Log on the MSDP server.
- 2 Use the following command to add the alternate interfaces:

```
cacontrol --interface add msdp-a.server.com
```

You can remove an added interface using the following command:

```
cacontrol --interface remove msdp-a.server.com
```

- 3 Use either of the following options to validate the interface configuration:
 - `cacontrol --interface list`
 - `bpstsinfo -si -storage_server msdp-a.server.com -stype PureDisk`
 Location of the `bpstsinfo` command:
`/usr/openv/netbackup/bin/admincmd/`
- 4 Configure the NetBackup client-side deduplication backup policy and run the backup operation.

About MSDP multi-domain support

An MSDP storage server is configured in a NetBackup media server. The NetBackup media servers and clients in the NetBackup domain can use this storage server. By default, the NetBackup media servers and clients cannot directly use an MSDP storage server from another NetBackup domain. For example, NetBackup media servers or clients cannot backup data to an MSDP storage server from another NetBackup domain.

To use an MSDP storage server from another NetBackup domain, the MSDP storage server must have multiple MSDP users. Then NetBackup media servers or clients can use the MSDP storage server from another NetBackup domain by using a different MSDP user. Multiple NetBackup domains can use the same MSDP storage server, but each NetBackup domain must use a different MSDP user to access that MSDP storage server.

To add an MSDP user on an MSDP storage server, run the following command:

- Windows

```
<install_path>\pdde\spauser -a -u <username> -p <password>
```

- UNIX

```
/usr/opensv/pdde/pdcr/bin/spauser -a -u <username> -p <password>
```

To list all the MSDP users, run the following command on the MSDP storage server:

- Windows

```
<install_path>\pdde\spauser -l
```

- UNIX

```
/usr/opensv/pdde/pdcr/bin/spauser -l
```

To use an MSDP storage server from another NetBackup domain, you must obtain a NetBackup certificate from another NetBackup domain.

Run the following commands on every NetBackup media server or client that wants to use an MSDP storage server from another domain:

- Windows

```
install_path\NetBackup\bin\NBCertcmd -getCACertificate -server  
another_master_server  
install_path\NetBackup\bin\NBCertcmd -getCertificate -server  
another_master_server -token token_string
```

- UNIX

```
/usr/opensv/netbackup/bin/nbcertcmd -getCACertificate -server  
another_master_server
```

```
/usr/opensv/netbackup/bin/nbcertcmd -getCertificate -server
another_master_server -token token_string
```

Use either of these two methods to obtain the authorization tokens:

- **NetBackup Administration Console**
 - Log on the target NetBackup master server and open **Security Management > Certificate Management > Token Management**.
 - Click **Create Token** or right-click in the blank area of the **Token records** list view and select **New Token** to create a token.
- **NetBackup Commands**
 - Use the `bpbnet` command to log on the target NetBackup master server.
 - Use the `nbcertcmd` command to get the authorization tokens.

For more information on the commands, refer to the *NetBackup Commands Reference Guide*.

An example for using an MSDP storage server from another NetBackup domain

The following table describes the hierarchy that is used in the example:

NetBackup domain A	NetBackup domain B
masterA	masterB
mediaA1	mediaB
mediaA2	
clientA	

`MasterA` is the host name of the master server of NetBackup domain A and the domain contains two media servers (`mediaA1`, `mediaA2`), and one client (`clientA`). `MasterB` is the host name of the master server of NetBackup domain B and the domain contains one media server (`mediaB`).

Using the following sample steps, create an MSDP storage server in domain B and let domain A use the MSDP storage server:

1. Create an MSDP storage server on media server `mediaB` of NetBackup domain B. (NetBackup Administration Console > **Media and Device Management > Configure Disk storage servers > Media Server Deduplication Pool**)
2. Run the following command on `mediaB` to create a new MSDP user `testuser1` with password as `testuser1pass`.

```
spausers -a -u "testuser1" -p "testuser1pass"
```

3. Run the following command on `mediaA1` to get a CA certificate and a host certificate from `masterB`.

```
nbcertcmd -GetCACertificate -server masterB
```

```
nbcertcmd -GetCertificate -server masterB -token <token_string>
```

4. Create an MSDP OpenStorage server on `mediaA1` of NetBackup domain A. (NetBackup Administration Console > **Media and Device Management** > **Configure Disk storage servers** > **OpenStorage**).

OpenStorage server type is **PureDisk**, Storage server name is `mediaB`, user name is `testuser1`, and password is `testuser1pass`.

You must enter the server type as **PureDisk**.

Now `mediaA1` of NetBackup domain can use the MSDP storage server `mediaB`. To use `mediaA2` as a load balance server of the MSDP storage server, you can run the following certificate command on `mediaA2`:

- `nbcertcmd -GetCACertificate -server masterB`

- `nbcertcmd -GetCertificate -server masterB -token <token_string>`

To run client-direct backup from `clientA` to MSDP storage server `mediaB`, run the following certificate command on `clientA`:

- `nbcertcmd -GetCACertificate -server masterB`

- `nbcertcmd -GetCertificate -server masterB -token <token_string>`

5. After creating the MSDP OpenStorage server, create a related NetBackup disk pool and storage unit. Use the storage unit to run all the related NetBackup jobs.

When optimized duplication and multi-domain are used together, there is communication between the MSDP storage servers from two different NetBackup domains. The MSDP storage server from the other domain must have a certificate generated by master server of the local NetBackup domain. Run the `nbcertcmd` commands on source side MSDP storage server to request a certificate from the NetBackup master server of target MSDP storage server.

When backup and restore jobs on the client and multi-domain are used together, there is communication between the NetBackup client and MSDP storage server from two different NetBackup domains. Run the `nbcertcmd` commands on the NetBackup client to request a certificate from the NetBackup master server of MSDP storage server.

When one NetBackup domain uses the MSDP storage server of another NetBackup domain, the MSDP storage server cannot be the A.I.R target of that NetBackup domain.

If an external CA is used in the NetBackup setup, you do not need to run the `nbcertcmd -GetCACertificate` and the `nbcertcmd -GetCertificate` commands. If NetBackup domains A and B do not use the same external CA, synchronize the external root CA between the two NetBackup domains for MSDP communication.

For more information of the external CA, refer to *NetBackup Security and Encryption Guide*.

When one NetBackup domain uses an MSDP storage server that has multiple network interfaces and related host names, another NetBackup domain can use any one host name to configure the OpenStorage server. If the MSDP storage server that has multiple host names uses an external CA, the **Subject Alternative Name** field of the external certificate must contain all the host names that are used to configure the OpenStorage server.

About MSDP mutli-domain VLAN Support

MSDP supports multi-domain NetBackup setups. In a multi-domain set-up, it is important for master servers from other domains to connect with the MSDP storage server and the master server of the NetBackup domain that contains the MSDP server. The master servers and media servers must have multiple network interfaces and host names in a multi-domain setup.

When you configure MSDP VLAN, the local NetBackup domain and the other NetBackup domain must have the NetBackup version 8.2 or later.

An example for using an MSDP VLAN

The following table describes the hierarchy that is used in the example:

NetBackup domain A	NetBackup domain B
masterA - (10.XX.30.1/24)	masterB - (10.XX.40.3/24)
masterA2 - (10.XX.40.1/24)	mediaB - (10.XX.40.4/24)
mediaA - (10.XX.30.2/24)	
mediaA2 - (10.XX.40.2/24)	

masterA is the master server of domain A and has two host names and IP addresses. mediaA is the media server of domain A and has two host names and IP addresses. MSDP storage server is created on media server mediaA.

To let domain B access the MSDP storage server on `mediaA` of domain A, run the following steps:

1. Create an MSDP storage server on media server `mediaA` of NetBackup domain A. (NetBackup Administration Console > **Media and Device Management** > **Configure Disk storage servers** > **Media Server Deduplication Pool**)
2. Run following command on `mediaA` to create a new MSDP user `testuser1` with password `testuser1pass`:

```
spauser -a -u "testuser1" -p "testuser1pass"
```

3. Servers in the domain B can only access IP like 10.XX.40.*, so `masterA2` is used as the master server host name of domain A.

Run following command on `mediaB` to get a CA certificate and a host certificate from `masterA`:

```
nbcertcmd -GetCACertificate -server masterA2
```

```
nbcertcmd -GetCertificate -server masterA2 -token <token_string>
```

If the `nbcertcmd -GetCACertificate` displays the error "The server name does not match any of the host names listed in the server's certificate", refer to the following article to add more host name to master server:

https://www.veritas.com/support/en_US/article.100034092

4. Create an MSDP OpenStorage server on `mediaB` of NetBackup domain B. NetBackup Administration Console > **Media and Device Management** > **Configure Disk storage servers** > **OpenStorage**).

The OpenStorage server name `mediaA2` is used as the host name that has the IP address 10.XX.40.*.

OpenStorage server type is **PureDisk**, user name is `testuser1`, and password is `testuser1pass`. You must enter the server type as **PureDisk**.

Now `mediaB` of NetBackup domain B can use the MSDP storage server `mediaA2` and the network IP address 10.XX.40.*

If an external CA is used in the NetBackup setup, you do not need to run the `nbcertcmd -GetCACertificate` and the `nbcertcmd -GetCertificate` commands. If NetBackup domain A and NetBackup domain B do not use the same external CA, you must synchronize the external root CA between the two NetBackup domains for MSDP communication. If the servers have multiple host names, then the **Subject Alternative Name** field of the external certificate must contain all the host names.

Configuring deduplication to the cloud with NetBackup Cloud Catalyst

This chapter includes the following topics:

- [Using NetBackup Cloud Catalyst to upload deduplicated data to the cloud](#)
- [Cloud Catalyst requirements and limitations](#)
- [Configuring a Linux media server as a Cloud Catalyst storage server](#)
- [Configuring a Cloud Catalyst storage server for deduplication to the cloud](#)
- [About the Cloud Catalyst esfs.json configuration file](#)
- [About the Cloud Catalyst cache](#)
- [Controlling data traffic to the cloud when using Cloud Catalyst](#)
- [Configuring source control or target control optimized duplication for Cloud Catalyst](#)
- [Configuring a Cloud Catalyst storage server as the source for optimized duplication](#)
- [Decommissioning Cloud Catalyst cloud storage](#)
- [NetBackup Cloud Catalyst workflow processes](#)
- [Disaster recovery for Cloud Catalyst](#)
- [About image sharing in cloud using Cloud Catalyst](#)

Using NetBackup Cloud Catalyst to upload deduplicated data to the cloud

NetBackup Cloud Catalyst uses MSDP deduplication technology to upload deduplicated data to the cloud. The data is uploaded by a Cloud Catalyst storage server, which first stores data in a local cache before uploading it to cloud storage. This cloud storage server is a dedicated host that can be either a NetBackup appliance or a media server configured for NetBackup Cloud Catalyst.

The following topics describe scenarios that use NetBackup Cloud Catalyst:

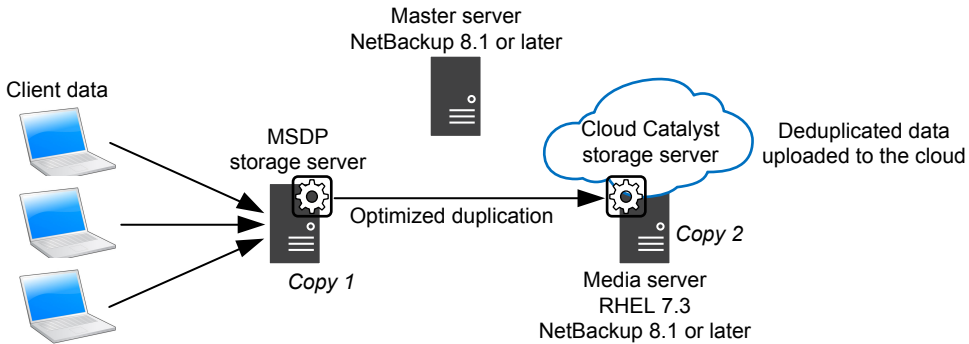
- See [“Optimized duplication is used to copy data from an MSDP storage server to a Cloud Catalyst storage server \(preferred use case\)”](#) on page 254.
This is the preferred use case because it uses the Cloud Catalyst storage server as a target for secondary copies. These copies usually have a longer retention and are uploaded to the cloud.
- See [“Backups go directly to a Cloud Catalyst storage server”](#) on page 256.
While this configuration is supported, it is not the preferred use case. In this use case, the Cloud Catalyst storage server is required to back up, deduplicate, and upload the data to cloud storage.

Optimized duplication is used to copy data from an MSDP storage server to a Cloud Catalyst storage server (preferred use case)

This is the preferred Cloud Catalyst use case. In this scenario, the NetBackup environment contains two media servers: one is an MSDP storage server and one is a Cloud Catalyst storage server.

- First, client data is backed up to the MSDP storage server, which is used for short-term data retention. In subsequent backups, the data is deduplicated. (See [Figure 6-1.](#))
- A storage lifecycle policy then copies the data to a Cloud Catalyst storage server using optimized duplication. Data is uploaded from the Cloud Catalyst local cache to the cloud when an MSDP container file is full. This occurs soon after the backup or duplication job begins, but not immediately.
The cloud storage is used for long-term data retention. Deduplication makes it possible for subsequent jobs to transfer substantially less data to the cloud, depending on the deduplication rate.

Figure 6-1 Flow of data from an MSDP storage server to a Cloud Catalyst storage server



This use case is preferred to backing up directly to a Cloud Catalyst storage server. (See [“Backups go directly to a Cloud Catalyst storage server”](#) on page 256.)


See [“Configuring a Cloud Catalyst storage server for deduplication to the cloud”](#) on page 261.

MSDP storage servers fan-in to a single Cloud Catalyst storage server

Multiple MSDP storage servers in a NetBackup domain can use a single Cloud Catalyst storage server, as shown in [Figure 6-2](#).

[Figure 6-2](#) also shows that multiple Cloud Catalyst storage servers can be included in one NetBackup domain. Separate Cloud Catalyst storage servers are required for each cloud provider that is used. One Cloud Catalyst storage server can write to only one cloud provider, and to only one cloud bucket.

Figure 6-2 Fan-in MSDP storage servers and multiple Cloud Catalyst storage servers



Backups go directly to a Cloud Catalyst storage server

In this scenario, the NetBackup environment does not use an MSDP storage server; it contains only a Cloud Catalyst storage server. The Cloud Catalyst storage server must back up, deduplicate, and upload the data to cloud storage.

This scenario is not the preferred Cloud Catalyst use case. This scenario may be a good solution for some users depending on the size of the environment, the desired backup window, and other factors. However, it is expected that most users get better results from the preferred use case.


See [“Optimized duplication is used to copy data from an MSDP storage server to a Cloud Catalyst storage server \(preferred use case\)”](#) on page 254.

Note the following disadvantages to backing up directly to the Cloud Catalyst storage server:

- Since the data needs to be uploaded to the cloud before a job can be considered complete, the job may not complete within the backup window.

- If a restore is required, the data may no longer be in the local cache. If the data has been uploaded to the cloud, the restore may be slow and expensive.
- Client Direct is not supported. To review the limitations, See [“Cloud Catalyst requirements and limitations”](#) on page 257.

Figure 6-3 Flow of data directly to a Cloud Catalyst storage server



Cloud Catalyst requirements and limitations

The Cloud Catalyst storage server is a dedicated media server for MSDP deduplicated cloud storage. The media server must be either a NetBackup appliance or a Linux media server, as described in [Table 6-1](#).

Table 6-1 The types of media servers that can be used as Cloud Catalyst storage servers

Host	Version
NetBackup appliance	Veritas NetBackup Cloud Catalyst Appliance
NetBackup media server	Red Hat Enterprise Linux, 7.3 or later NetBackup 8.1 or later See “Configuring a Linux media server as a Cloud Catalyst storage server” on page 259.

The Cloud Catalyst storage server can work with a proxy server when the NTLM authentication method is used. The proxy server must follow certain requirements to function inside of the NetBackup system. The following items are required only

when the NTLM authentication method is used. When the `BASIC` authentication method or `NONE` is selected, there is no such requirement.

Requirements for a proxy server:

- The proxy server must support both NTLMv1 and NTLMv2.
- Cloud Catalyst is dependent on Classic Cloud Connector to get the list of buckets.
 - The Classic Cloud Connector only supports NTLMv1, and not NTLMv2.
- Cloud Catalyst only supports NTLMv2.

Note: A Cloud Catalyst storage server is supported for Auto Image Replication (A.I.R.) and Target A.I.R. as the source and the target from NetBackup 8.3 or later.

Limitations

NetBackup Cloud Catalyst recognizes nearly all other MSDP features and capabilities.

See the NetBackup compatibility lists for updated information on supported cloud vendors and feature support:

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

The following are current exclusions or limitations:

- Support limitations as listed in the *NetBackup Cloud Administrator's Guide*:
<http://www.veritas.com/docs/DOC5332>
- The NetBackup media servers that you use as Cloud Catalyst storage servers must be the same NetBackup version as the master server.
- Load balancing for Cloud Catalyst storage servers or Cloud Catalyst appliances is supported. See the following topic for configuration considerations.
See “[Configuring source control or target control optimized duplication for Cloud Catalyst](#)” on page 277.
- OpsCenter cannot report on the data that Cloud Catalyst storage servers upload to the cloud.
- A single media server cannot be both an MSDP storage server and a Cloud Catalyst storage server.
- A Cloud Catalyst storage server is supported as the target for A.I.R. when the storage server is selected from the drop-down list of the trusted master server.
- NetBackup Client Direct deduplication is supported in a Cloud Catalyst duplication to cloud scenario. However, Client Direct deduplication is not supported in a

scenario where the backup is uploaded directly to the cloud and no duplication is configured.

- Direct backup to Cloud Catalyst using the accelerator option is not supported when Amazon S3 Glacier or Amazon S3 Glacier Deep Archive storage classes are used. Backup to an MSDP with accelerator and then optimized duplication to Cloud Catalyst is supported.
- NetBackup clients must be on version 8.3 or later to support Client Direct enabled backups to the Cloud Catalyst media server.

Configuring a Linux media server as a Cloud Catalyst storage server

The Cloud Catalyst storage server is a dedicated media server for MSDP deduplicated cloud storage.

The media server that is configured as the Cloud Catalyst storage server must be one of the following host types:

- NetBackup Cloud Catalyst Appliance
Documentation for the Cloud Catalyst Appliance can be found at the following website:
https://www.veritas.com/content/support/en_US/Appliances.html
- A Linux media server with specifications similar to the NetBackup Cloud Catalyst Appliance. The following section lists the software and hardware requirements for this media server.

Specifications for a Linux media server

The dedicated media server that will be configured as a Cloud Catalyst storage server should meet or exceed the specifications of the NetBackup Cloud Catalyst Appliance. The requirements for this MSDP media server are greater than the minimum MSDP server requirements:

See “[About MSDP server requirements](#)” on page 28.

In addition to the technical requirements that are listed in [Table 6-2](#), the media server must also contain the following:

- Red Hat Enterprise Linux 7.3 or later.
- NetBackup 8.1 or later.
- A minimum of 4 TB available for the Cloud Catalyst local cache directory.
See “[About the Cloud Catalyst cache](#)” on page 275.
- The processor must support the SSE4.2 instruction set.

The technical specifications for the NetBackup 5240 Appliance are listed in the appendix of the *NetBackup 5240 Appliance Product Description*. [Table 6-2](#) lists a subset of minimum requirements for the media server.

Table 6-2 NetBackup 5240 Appliance technical specifications

Technical Specification	NetBackup 5240 Appliance
Processor	Two Intel® Xenon® E5-2630 v3 2.40 GHz processors
CPU speed	2.40 GHz
Cores	16 (8 per processor)
System memory	192 GB DDR4 RDIMM
Usable MSDP storage capacity	14 TB with 1 GB RAID 6 cache Note: If a Linux media server (non-appliance) is used, it must have a minimum of 4 TB available for the Cloud Catalyst local cache directory. See “About the Cloud Catalyst cache” on page 275.
10 Gb Ethernet ports	5

If you don't use a NetBackup appliance, the following table describes the minimum system requirements for a media server that is configured for Cloud Catalyst.

Table 6-3 NetBackup minimum system requirements

Hardware	Minimum specifications
CPU speed	2.4 GHz
Cores	4
System memory	16 GB
10 Gb Ethernet ports	1 Note: Multiple networking interface cards (NIC) may be desired for redundancy and connectivity reasons.
Cache size	1 TB
CPU instruction set support	SSE 4.2
Aggregate throughput	1,000 MB/s

Table 6-3 NetBackup minimum system requirements (continued)

Hardware	Minimum specifications
Optimal streams	16

Configuring a Cloud Catalyst storage server for deduplication to the cloud

The configuration of deduplication to the cloud differs, depending on the type of host that is used for the Cloud Catalyst storage server.

- See the following procedure if you use a NetBackup Cloud Catalyst Appliance:
See [“How to configure a NetBackup Cloud Catalyst Appliance”](#) on page 261.
- See the following procedure if you use a Linux media server (non-appliance):
See [“How to configure a Linux media server as a Cloud Catalyst storage server”](#) on page 265.

Note: Cloud Catalyst uses cloud connectors during initial configuration. Certificate validation against Certificate Revocation List (CRL) secures the calls that the cloud connectors initiate. After configuration, `nbrmms` periodically polls Cloud Catalyst storage server and makes a call to the cloud connector to check its state using the cloud connector’s CRL.

For more information on Certificate Revocation List (CRL), Review the section *Certificate validation against Certificate Revocation List (CRL)* in [NetBackup Cloud Administrator’s Guide](#).

After initial configuration, Cloud Catalyst uses Online Certificate Status Protocol (OCSP) for securing calls during data transfer to and from the cloud. Review the following section for more information on OCSP:

See [“Certificate validation using Online Certificate Status Protocol \(OCSP\)”](#) on page 269.

How to configure a NetBackup Cloud Catalyst Appliance

Use the following procedure to configure a NetBackup Cloud Catalyst Appliance as a Cloud Catalyst storage server.

This procedure assumes that the preferred use case is configured. The preferred use case is described in the following topic:

See [“Optimized duplication is used to copy data from an MSDP storage server to a Cloud Catalyst storage server \(preferred use case\)”](#) on page 254.

Table 6-4 Configuration steps for a NetBackup Cloud Catalyst Appliance

Step	Task	Where to configure	Procedure
Step 1	Configure an MSDP storage server (source)	One or more media servers	<p>A NetBackup environment may contain one or more media servers that serve as source MSDP storage servers for one or more Cloud Catalyst appliances.</p> <p>See “About MSDP optimized duplication within the same domain” on page 141.</p> <p>See “About the media servers for MSDP optimized duplication within the same domain” on page 144.</p>
Step 2	Create an account with a cloud service provider	Cloud service provider	<p>Make sure that you have an account with a cloud service provider and that you know the user name and password.</p>
Step 3	Configure KMS (optional)	Master server	<p>NetBackup Cloud Catalyst supports MSDP Key Management Service (KMS) encryption.</p> <p>NetBackup supports NetBackup KMS and external KMS for data encryption.</p> <p>To configure KMS and encrypt backups, log into the appliance as a NetBackup CLI user to run the commands.</p> <p>For more information about KMS configuration, see the NetBackup Security and Encryption Guide.</p>

Table 6-4 Configuration steps for a NetBackup Cloud Catalyst Appliance
(continued)

Step	Task	Where to configure	Procedure
Step 4	Configure an appliance as a Cloud Catalyst storage server (target)	Appliance	<p>A Cloud Catalyst Appliance is the target for the optimized duplication. It is a dedicated appliance for cloud storage.</p> <p>A NetBackup environment may contain one or more Cloud Catalyst appliances. There should be one appliance for each different cloud service provider that is used in the environment.</p> <p>Perform the initial configuration on NetBackup Appliance Shell Menu, then run the NetBackup Appliance Cloud Storage Server Configuration Wizard in the NetBackup Administration Console.</p> <p>If desired, select the Encrypt data using MSDP KMS encryption in the wizard to configure encryption.</p> <p>If the Encrypt data using MSDP KMS encryption option is selected, one of following operations is carried out:</p> <ul style="list-style-type: none"> ■ If none of the KMS servers is configured, by default NetBackup KMS is configured. ■ If NetBackup KMS server is configured, you need to provide key name and key passphrase. ■ If external KMS server is configured on the master server, you must enter key group name. You need to ensure that a key with a custom attribute set with value of key group name already exists in the external KMS server. <p>For more information about KMS configuration, see the NetBackup Security and Encryption Guide.</p> <p>As part of this wizard, a disk pool and a storage unit are configured.</p> <p>The default path for the Local cache directory: is <code>/msdpc</code>. You can change this directory by clicking on Browse.</p> <p>Documentation for the Cloud Catalyst Appliance can be found at the following website:</p> <p>https://www.veritas.com/content/support/en_US/Appliances.html</p>

Table 6-4 Configuration steps for a NetBackup Cloud Catalyst Appliance
(continued)

Step	Task	Where to configure	Procedure
Step 5	Configure the target to allow load balancing for push optimized duplications (optional)	Appliance	<p>To configure load balancing for push optimized duplications, the Cloud Catalyst Appliance must allow access from the source media server.</p> <p>Note: The push configuration is recommended if using load balancing for Cloud Catalyst.</p> <p>See the following topic for configuration considerations.</p> <p>See “Configuring source control or target control optimized duplication for Cloud Catalyst” on page 277.</p>
Step 6	Configure a storage lifecycle policy	Master server	<p>See “Configuring a storage lifecycle policy for NetBackup Cloud Catalyst” on page 270.</p> <p>Note: The backup and the duplication jobs can also be run manually from the NetBackup Administration Console or by using the <code>bpduplicate</code> command.</p>
Step 7	Configure a backup policy	Master server	<p>Create a backup policy that uses the SLP. In the Policy storage field of the policy Attributes tab, select the SLP.</p> <p>See “About MSDP backup policy configuration” on page 199.</p> <p>See “Creating a backup policy” on page 200.</p>
Step 8	Protect the MSDP data and catalog	Media server	<ul style="list-style-type: none"> Manually create an MSDP catalog backup policy for every MSDP storage server and Cloud Catalyst storage server in the environment. <p>See “About protecting the MSDP catalog” on page 233.</p> <p>See “Configuring an MSDP catalog backup” on page 239.</p> <p>Note: After the MSDP catalog backup policy is created, it is inactive. Assign a storage unit to it and then make it active by setting the Go into effect at option on the Attributes tab.</p> Save a configuration file for every MSDP storage server and Cloud Catalyst storage server in the environment. <p>See “About saving the MSDP storage server configuration” on page 227.</p>

How to configure a Linux media server as a Cloud Catalyst storage server

Use the following procedure if you use a Linux media server (non-appliance) as a Cloud Catalyst storage server.

This procedure assumes that the preferred use case is configured. The preferred use case is described in the following topic:

See [“Optimized duplication is used to copy data from an MSDP storage server to a Cloud Catalyst storage server \(preferred use case\)”](#) on page 254.

Table 6-5 Configuration steps for a media server (non-appliance)

Step	Task	Where to configure	Procedure
Step 1	Configure an MSDP storage server (source)	One or more media servers	<p>A NetBackup environment may contain one or more media servers that serve as source MSDP storage servers for Cloud Catalyst storage servers.</p> <p>See “About MSDP optimized duplication within the same domain” on page 141.</p> <p>See “About the media servers for MSDP optimized duplication within the same domain” on page 144.</p>
Step 2	Create an account with a cloud service provider	Cloud service provider	Make sure that you have an account with a cloud service provider and that you know the user name and password.
Step 3 (Conditional)	Deploy host name-based certificates	Master server	<p>This step is applicable and required for media server versions 8.0 to 8.1.2 only.</p> <p>Deploy a host name-based security certificate for the media server that is configured as the Cloud Catalyst storage server.</p> <ol style="list-style-type: none"> <code>/usr/opensv/netbackup/bin/admincmd/bpnbaz -ProvisionCert host_name</code> Restart the services on the media server after generating a certificate.

Table 6-5 Configuration steps for a media server (non-appliance)
(continued)

Step	Task	Where to configure	Procedure
Step 4	Configure a Cloud Catalyst storage server (target)	Media server	<p>A Cloud Catalyst storage server is the target for the optimized duplication. It is a dedicated media server for cloud storage.</p> <ul style="list-style-type: none"> Run the Cloud Storage Server Configuration Wizard in the Java GUI selecting the Enable NetBackup Cloud Catalyst check box and entering a path for the local cache directory. The administrator configures a local cache directory as part of configuring a Cloud Catalyst storage server. The minimum size for this directory is 4 TB, however, an environment may see benefits from a larger local cache directory. See “About the Cloud Catalyst cache” on page 275. If desired, select the Encrypt data using MSDP KMS encryption in the wizard to configure encryption. If the Encrypt data using MSDP KMS encryption option is selected, one of following operations is carried out: <ul style="list-style-type: none"> If none of the KMS servers is configured, by default NetBackup KMS is configured. If NetBackup KMS server is configured, you need to provide key name and key passphrase. If external KMS server is configured on the master server, you must enter key group name. You need to ensure that a key with a custom attribute set with value of key group name already exists in the external KMS server. For more information about KMS configuration, see the NetBackup Security and Encryption Guide. A NetBackup environment may contain one or more media servers that serve as Cloud Catalyst storage servers. There should be one Cloud Catalyst storage server for each different cloud service provider that is used in the environment.
Step 5	Ensure that a common media server is configured to allow optimized duplication	Media server	<p>If the Cloud Catalyst media server was added as a load balancing media server during the MSDP storage server wizard (Step 1), pull optimized duplication is configured.</p> <p>See the following topic for configuration considerations: See “Configuring source control or target control optimized duplication for Cloud Catalyst” on page 277.</p>

Table 6-5 Configuration steps for a media server (non-appliance)
(continued)

Step	Task	Where to configure	Procedure
Step 6	Protect the MSDP data and catalog	Media server	<ul style="list-style-type: none"> Create an MSDP catalog backup policy for every MSDP storage server and Cloud Catalyst storage server in the environment. See “About protecting the MSDP catalog” on page 233. See “Configuring an MSDP catalog backup” on page 239. Note: After the MSDP catalog backup policy is created, it is inactive. Assign a storage unit to it and then make it active by setting the Go into effect at option on the Attributes tab. Save a configuration file for every MSDP storage server and Cloud Catalyst storage server in the environment. See “About saving the MSDP storage server configuration” on page 227.
Step 7	Configure a storage lifecycle policy	Master server	<p>See “Configuring a storage lifecycle policy for NetBackup Cloud Catalyst” on page 270.</p> <p>The backup and the duplication jobs can also be run manually from the NetBackup Administration Console or by using the <code>bpduplicate</code> command.</p>
Step 8	Configure a backup policy	Master server	<p>Create a backup policy that uses the SLP. In the Policy storage field of the policy Attributes tab, select the SLP.</p> <p>See “About MSDP backup policy configuration” on page 199.</p> <p>See “Creating a backup policy” on page 200.</p>

Configuring a Cloud Catalyst storage server as the target for the deduplications from MSDP storage servers

This media server is the target for the backups from the MSDP storage server. After configuration, this storage server is the Cloud Catalyst storage server that uploads the deduplicated data to the cloud. The following describes a traditional NetBackup server configuration. The NetBackup Appliance Shell Menu has similar steps for configuration.

Configuring a Cloud Catalyst storage server as the target for the MSDP storage servers

- 1** On the master server, in the **NetBackup Administration Console**, click **Configure Cloud Storage Servers** to launch the Cloud Storage Server Configuration Wizard.
- 2** In the wizard, enable the option: **Enable NetBackup Cloud Catalyst**.
- 3** Enter a path to use as the **Local cache directory**. Deduplicated data is written to this directory on the cloud storage server before it is uploaded to the cloud.

Notes:

- The minimum size for the local cache directory is 4 TB, however, an environment may see benefits from a larger local cache directory.
- This path should be to a file system which is dedicated for Cloud Catalyst cache use. Inaccurate cache eviction occurs if the path shares any storage with other data or applications.
- NetBackup manages the files in the local cache directory. Users should not manually delete files in this directory.
- The *NetBackup Cloud Administrator's Guide* contains information about all of the available options and switches which are common to non-Cloud Catalyst and Cloud Catalyst environments.
<http://www.veritas.com/docs/DOC5332>
 See "About the Cloud Catalyst cache" on page 275.

- 4** Choose whether to use NetBackup KMS encryption. To enable encryption, select the **Encrypt data using MSDP KMS encryption** in the wizard to configure encryption.

If the **Encrypt data using MSDP KMS encryption** option is selected, one of following operations is carried out:

- If none of the KMS servers is configured, by default NetBackup KMS is configured.
- If NetBackup KMS server is configured, you need to provide key name and key passphrase.
- If external KMS server is configured on the master server, you must enter key group name. You need to ensure that a key with a custom attribute set with value of key group name already exists in the external KMS server.

For more information about KMS configuration, see the [NetBackup Security and Encryption Guide](#).

- 5 As part of the storage server wizard, a disk pool and a storage unit are configured.

When the storage lifecycle policy is created, this storage unit name is used in the **Storage unit** field for the **Duplication** operation.

See “[Configuring a storage lifecycle policy for NetBackup Cloud Catalyst](#)” on page 270.

Certificate validation using Online Certificate Status Protocol (OCSP)

For all the cloud providers, Cloud Catalyst provides a capability to verify the SSL certificates. Cloud Catalyst uses a combination of Certificate Revocation List (CRL) and Online Certificate Status Protocol (OCSP). During initial configuration, Cloud Catalyst cloud connectors perform certificate validation using CRL. After initial configuration, Cloud Catalyst uses OCSP for securing calls during data transfer to and from the cloud. OCSP can be enabled along with CRL and if SSL and CRL options are enabled, each externally signed SSL certificate is verified using the OCSP. If the certificate is revoked, Cloud Catalyst does not connect to the cloud provider.

You can enable validation using OCSP in the same way as CRL. Refer to *Certificate validation against Certificate Revocation List (CRL)* in [NetBackup Cloud Administrator's Guide](#) for more details.

Requirements for enabling the certificate validation using OCSP:

- OCSP responder endpoints are `HTTP`. Disable any firewall rule that blocks `HTTP` (port 80) connection to external network. For example, <http://ocsp.msocsp.com>.
- OCSP responder URL is dynamically fetched from the certificate. Disable any firewall rule that blocks unknown URLs.
- OCSP responder URL must be present in the x.509 certificate. The type of OCSP responder URL must be `HTTP`.
- Private Clouds typically have a self-signed certificate and don't need certificate verification. The OCSP check is skipped in this case, regardless of whether CRL or OCSP is enabled or not.

Note: When a storage server has SSL enabled, the OCSP validation is enabled by default when you upgrade to NetBackup Cloud Catalyst 8.2.1.

Managing Cloud Catalyst storage server with IAM Role or CREDS_CAPS credential broker type

When you use the **NetBackup Administration Console** to configure Cloud Catalyst storage server using IAM Role or CREDS_CAPS credential broker type, the default credentials are set as user name: **admin** and password: **password**.

You can update these credentials using the `-tpconfig` command. Use the same command, to configure the Cloud Catalyst storage server using the command line.

Example:

```
tpconfig -add -storage_server storage_server_name -stype  
PureDisk_amazon_rawd -sts_user_id "my-username" -password  
"my-password"
```

Note: Remember these credentials because the same credentials are required during disaster recovery of the Cloud Catalyst storage server.

Configuring a storage lifecycle policy for NetBackup Cloud Catalyst

A storage lifecycle policy (SLP) manages the backup jobs and the duplication jobs from a MSDP storage server to a MSDP cloud storage server. The backup and duplication jobs can also be run manually from the **NetBackup Administration Console** or by using the `bpduplicate` command.

This procedure assumes that the preferred use case is being configured. The preferred use case is described in the following topic:

See [“Optimized duplication is used to copy data from an MSDP storage server to a Cloud Catalyst storage server \(preferred use case\)”](#) on page 254.

If an SLP is used, it must contain a **Backup** operation and a child **Duplication** operation.

Configuring a storage lifecycle policy for deduplication to the cloud

- 1 The first operation must be a **Backup** operation. In the **Storage unit** field, select the storage unit that was created for the MSDP storage server.
- 2 Create a second operation that is a child of the **Backup** operation. The second operation must be a **Duplication** operation. In the **Storage unit** field, select the storage unit that was created for the Cloud Catalyst storage server.

See [“Configuring a Cloud Catalyst storage server as the target for the deduplications from MSDP storage servers”](#) on page 267.

About the Cloud Catalyst `esfs.json` configuration file

NetBackup Cloud Catalyst uses the configuration options that are included in the `esfs.json` file for many operations. They include the options that determine when the files that have been uploaded to cloud storage are removed (evicted) from the local cache. Some of the options are configurable by the NetBackup administrator. The location of the `esfs.json` file depends on the location of the local cache directory.

As part of the Cloud Storage Server Configuration Wizard, the administrator configures a local cache directory. The local cache directory (`local_cache_dir` in the following topic) determines the location of other directories which are installed automatically and are critical to Cloud Catalyst operations:

- A mount path (`/storage`): `local_cache_dir/storage`
The Veritas NetBackup Extendable Storage File System process (`vxesfsd`) runs on the mount path.
- A cache path (`/cache`): `local_cache_dir/cache/etc/esfs.json`
This directory is the location of the `esfs.json` file.

Configuring the options in the `esfs.json` file

The NetBackup Cloud Catalyst configuration options are listed in [Table 6-6](#).

To change the configuration items in the `esfs.json` configuration

- 1 Change one or more of the configuration items in the `esfs.json` file. The file is found in the following location:

```
local_cache_dir/cache/etc/esfs.json
```

- 2 Save and close the file.
- 3 If `vxesfsd` is running, run the `esfs_reconfig` command (enter the mount path) as follows:

```
/usr/opensv/esfs/bin/esfs_reconfig local_cache_dir/storage
```

If `esfs.json` is changed while `vxesfsd` is stopped, the changes take effect the next time `vxesfsd` starts.

- 4 Some items require that you restart `vxesfsd` before the new configuration can take effect.

These items are indicated in [Table 6-6](#).

- Before `vxesfsd` is stopped, verify that no processes currently use `vxesfsd`, including the current working directory of any user.

The NetBackup Deduplication Manager (*spad*) and the NetBackup Deduplication Engine (*spoold*) use *vxesfsd*, so they need to be stopped if either is running.

- Before *vxesfsd* is started, make sure that no data exists in the mount point. If data exists in the mount point, *vxesfsd* fails to start. See [“Restarting the vxesfsd process”](#) on page 441.

Table 6-6 Configuration items in the esfs.json file

Configuration item	Default setting	Description	<i>vxesfsd</i> restart required?
Log			
Path	/usr/opensv/ netbackup/logs	<p>This setting indicates the directory where logs for the <i>vxesfsd</i> process are created.</p> <p>The logs include the following:</p> <ul style="list-style-type: none"> ■ <i>esfs_database</i> ■ <i>esfs_filesystem</i> ■ <i>esfs_monitor</i> ■ <i>esfs_storage</i> ■ <i>esfs_storagemanager</i> 	Yes
Size	10485760	<p>This value (in bytes) controls the maximum size that a single log file is allowed to grow. Once the file reaches that approximate size, the log file is closed and another log file is opened.</p>	No
Level	1	<p>This value determines the logging level and what information is included in the logs:</p> <ul style="list-style-type: none"> ■ 1: Errors only ■ 2: Warnings and errors ■ 3: All log entries <p>Note: The Global logging level in the Host Properties has no effect on this setting.</p>	No. Changes may not take effect until the next restart.
Monitor			

Table 6-6 Configuration items in the esfs.json file (*continued*)

Configuration item	Default setting	Description	vxesfsd restart required?
HighWatermark	90	<p>This value controls how full the Cloud Catalyst cache partition can become. The value represents the percentage of used space on the cache partition. If the <code>HighWatermark</code> is hit, Cloud Catalyst suspends writes to the cache. Cloud Catalyst operations (for example backup and duplication operations to the Cloud Catalyst server) are paused until space is made available through cache eviction. Together with the <code>LowWatermark</code>, this value also determines when cache eviction starts</p> <p>Cache eviction starts when the used space reaches the midpoint between <code>HighWatermark</code> and <code>LowWatermark</code> $(high+low)/2$ and continues until <code>LowWatermark</code> is reached.</p> <p>If the <code>LowWatermark</code> value is set very low, the cache is nearly cleared. That is, nearly all of the eligible files are evicted.</p> <p>Note: The file system on which the cache resides must not contain data from any other applications.</p>	No
LowWatermark	50	<p>This value determines at what point <code>vxesfsd</code> stops evicting container files after the <code>HighWatermark</code> triggers eviction. The value represents the percentage of used space on the cache partition.</p>	No
BackupDBTime	12	<p>This value determines how frequently (in hours) the database is backed up.</p> <p>This value is unrelated to the <code>drcontrol</code> policy, if one exists. Veritas recommends that a <code>drcontrol</code> policy is created and used to protect your cloud storage server. Veritas recommends that this value is not changed.</p> <p>Note: This database is a special database for the cloud storage server, not the NetBackup or MSDP catalog.</p>	No

Table 6-6 Configuration items in the esfs.json file (*continued*)

Configuration item	Default setting	Description	vxesfsd restart required?
AsyncMaxConcurrentConnections	40	This value determines the maximum number of concurrent connections to the cloud provider for asynchronous (upload and delete) operations.	Yes
SyncMaxConcurrentConnections	40	This value determines the maximum number of concurrent connections to the cloud provider for synchronous (download) operations.	Yes
BackupDBCopies	14	<p>This value determines the maximum number of saved copies of the database that are created for backup.</p> <p>Veritas recommends that this value is not changed.</p> <p>This value is unrelated to the <code>drcontrol</code> policy, if one exists. Veritas recommends that a <code>drcontrol</code> policy is created and used to protect your cloud storage server.</p> <p>Note: This database is a special database for the cloud storage server, not the NetBackup or MSDP catalog.</p>	No
FileSystem			
MaxOpenFile	192114	<p>This value determines the maximum number of files that can be open at one time.</p> <p>Veritas recommends that this value is not changed.</p>	Yes
ReadOnly	0	<p>This value indicates whether the files currently stored in the cloud can be modified:</p> <ul style="list-style-type: none"> ■ 0: The files can be modified. (Default.) ■ 1: The files are read only. <p>This setting is used for disaster recovery and should not be changed under normal usage.</p>	Yes
ImageSharingEnabled	1	This value is used for disaster recovery in the Cloud. This option is enabled by default.	Yes

About the Cloud Catalyst cache

The administrator configures a local cache directory as part of configuring a Cloud Catalyst storage server. The primary function of the local cache directory (or Cloud Catalyst cache) is to allow the Cloud Catalyst to continue to deduplicate data. Deduplication of data occurs even if the ingest rate from targeted backup and duplication jobs temporarily exceeds the available upload throughput to the destination cloud storage.

For example, if backup and duplication jobs transfer 10 TB of data per hour to the Cloud Catalyst storage server, and the Cloud Catalyst deduplicates the data at a ratio of 10:1, the 1 TB of deduplicated data may exceed the upload capacity of .7 TB per hour of writes to cloud storage. The cache allows the jobs to continue to send and process the data, assuming that at some point the incoming data rate slows. The Cloud Catalyst cache only stores the deduplicated data. Jobs are not marked as complete until all data is uploaded to the cloud.

While a Cloud Catalyst cache of 4 TB is recommended, a larger cache has the following benefits:

- For restores:
If the data exists in the Cloud Catalyst cache, it is restored from the cache instead of the cloud. The larger the cache, the more deduplicated objects can reside in the cache.
- For data with poor deduplication rates:
A larger cache may be required since the poor deduplication ratios require that larger amounts of data be uploaded to the cloud.
- For job windows that experience bursts of activity:
A larger cache can be helpful if frequent jobs are targeted to the Cloud Catalyst storage server within a narrow window of time.

While a larger cache can be beneficial, jobs are not marked as complete until all data is uploaded to the cloud. Data is uploaded from the cache to the cloud when an MSDP container file is full. This occurs soon after the backup or duplication job begins, but not immediately. Deduplication makes it possible for second and subsequent backup jobs to transfer substantially less data to the cloud, depending on the deduplication rate.

See [“Optimized duplication is used to copy data from an MSDP storage server to a Cloud Catalyst storage server \(preferred use case\)”](#) on page 254.

For example, 4 TB of cache is expected to manage 1 PB of data in the cloud without issue.

Note: If you initiate a restore from Glacier or Glacier Deep Archive, NetBackup initiates a warming step. NetBackup does not proceed with the restore until all the data is available in S3 storage to be read.

The warming step is always done if using Amazon, even if the data is in the Cloud Catalyst cache. For storage classes other than Glacier and Glacier Deep Archive, the warming step is almost immediate with no meaningful delay. For Glacier and Glacier Deep Archive, the warming step may be immediate if files were previously warmed and are still in S3 Standard storage. However, it may take several minutes, hours, or days depending on settings being used.

How Cloud Catalyst manages the cache

The Cloud Catalyst manages the cache based on the configuration settings in the `esfs.json` file. Once the high watermark is reached, data is purged when the used space reaches the midpoint between `HighWatermark` and `LowWatermark` $(\text{high} + \text{low}) / 2$ and continues until `LowWatermark` is reached. If the rate of incoming data exceeds the rate where the watermark can be maintained, the jobs begin to fail. Administrators should not manually delete or purge the managed data in the cache storage unless directed to do so by NetBackup Technical Support.

See [“About the Cloud Catalyst esfs.json configuration file”](#) on page 271.

Controlling data traffic to the cloud when using Cloud Catalyst

Cloud Catalyst offers a throttling configuration option to limit the internet bandwidth that is used for uploading and downloading cloud data. Use this topic to enable throttling.

Note: By default, throttling is not enabled for NetBackup storage servers.

Change the Scalable Storage media server host properties:

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management** > **Host Properties** > **Media Servers**.
- 2 In the details pane, select the media server that is the Cloud Catalyst storage server.
- 3 On the **Actions** menu, select **Properties**.
- 4 Select the **Scalable Storage** properties.
- 5 Set **Total available bandwidth** to the desired bandwidth limit.

6 Set the **Sampling interval** to one or more to enable throttling (zero means that throttling is disabled).

7 Click **Apply** and **OK** to close the dialog.

See [“About the Cloud Catalyst esfs.json configuration file”](#) on page 271.

Configuring source control or target control optimized duplication for Cloud Catalyst

Optimized duplication requires the source storage server and the destination storage server to have at least one media server in common. For example, in a NetBackup Cloud Catalyst environment that includes an MSDP media server (the source storage server) and one Cloud Catalyst storage server (destination/target), either of these hosts could be the shared media server to perform optimized duplication.

The shared media server requires credentials for both the source storage server and the target storage server. (Credentials are granted in the **Media Servers** tab of the storage server properties for each storage server.) This point is important in determining whether your environment can support a source control or a target control operation.

The media server which manages the duplication operation determines whether the operation is a source controlled or a target controlled operation, as follows:

- If the MSDP media server manages the duplication operation, it is source controlled duplication. The Cloud Catalyst storage server must allow access from the MSDP media server. (See [Example source controlled configuration: Source media server must be RHEL and NetBackup 8.1 or later.](#))
- The Cloud Catalyst can allow access only if the MSDP media server is RHEL (any version) and NetBackup 8.1 or later.
 - If the MSDP media server is non-RHEL, it appears as a selection, but upon applying the credential, it fails with an **Add credentials failed** message.
 - If the MSDP media server is non-8.1, it does not appear in the **Media Servers** tab.
- If the MSDP media server is not RHEL or is not at NetBackup 8.1 or later, a target control operation must be used instead.
- If the Cloud Catalyst media server manages the duplication operation, it is target controlled duplication. The MSDP media server must allow access from the Cloud Catalyst media server. (See [Example target controlled configuration: Source media storage server is not RHEL or is pre-8.1.](#))

The MSDP media server can be of any supported platform or NetBackup version. See the considerations that are listed in [Table 6-7](#).

Table 6-7 Considerations for using source controlled or target controlled optimized duplication for Cloud Catalyst

Control of optimized duplication	Advantages and disadvantages
Source controlled	<p>Advantage: The NetBackup Disk Manager (<code>bpdm</code>) runs on the source media server, which prevents overloading the Cloud Catalyst storage server during duplication jobs.</p> <p>Disadvantage: Requires that the MSDP storage server be of a specific platform and NetBackup version.</p>
Target controlled	<p>Advantage: The MSDP media server can be of any supported platform or NetBackup version.</p> <p>Disadvantage: The NetBackup tape manager (<code>bptm</code>) may run on the Cloud Catalyst media server during backups. NetBackup may inadvertently use the Cloud Catalyst storage server for load balancing during backups.</p>

Note: In MSDP optimized duplication, the `spad` process on the source pool reads the source side data and sends the data to the target pool.

Note about credentials on media servers

There must be at least one media server which has credentials to both the MSDP storage server and the Cloud Catalyst storage server for jobs to be successful.

- If there is no media server which has credentials to both storage servers, jobs fail.
- If the common media server is the same host as the source MSDP storage server, it's a source control operation. ([Example source controlled configuration: Source media server must be RHEL and NetBackup 8.1 or later](#))
- If the common media server is the same host as the Cloud Catalyst storage server, it's a target control operation. ([Example target controlled configuration: Source media storage server is not RHEL or is pre-8.1](#))
- If the media servers for both have credentials in common, the NetBackup Resource Broker (NBRB) chooses a media server at job time. Some jobs can be a source control operation and some can be a target control operation.


Example source controlled configuration: Source media server must be RHEL and NetBackup 8.1 or later

In [Figure 6-4](#), the source media server (the MSDP storage server) is configured for a source control duplication to the target. Note that the source is Red Hat Enterprise Linux and has NetBackup 8.1 installed.

Two configuration items allow for source controlled duplication:

- The target (the Cloud Catalyst storage server) allows access from the source media server. The credential is granted in the storage server properties of each.
- The storage unit settings on the source and the target specifically point to the source media server. (Enable **Only use the following media servers** and select the source media server.)

Figure 6-4 Source controlled configuration



Storage server properties on source

This screenshot shows the 'Change Storage Server' dialog for the source server, 'MSDP_st_server'. The 'Server type' is 'PureDisk'. The 'Media Servers' tab is selected. The text explains that the deduplication workload can be load balanced to other media servers. Below, there is a list of media servers: 'Media Server' (unchecked), 'CCat_st_server' (unchecked), and 'MSDP_st_server' (checked).

Storage server properties on target

This screenshot shows the 'Change Storage Server' dialog for the target server, 'CCat_st_server'. The 'Server type' is 'PureDisk'. The 'Media Servers' tab is selected. The text explains that the deduplication workload can be load balanced to other media servers. Below, there is a list of media servers: 'Media Server' (unchecked), 'CCat_st_server' (checked), and 'MSDP_st_server' (checked).

Storage unit properties on source

This screenshot shows the 'New Storage Unit' dialog for the source storage unit, 'STU_source'. The 'Storage unit type' is 'Disk'. The 'Disk type' is 'PureDisk'. Under 'Properties and Server Selection', the 'Select disk pool' is 'DP_MSDP_source'. The 'Media server' section has 'Only use the following media servers' selected. The 'Media Servers' list shows 'CCat_st_server' (unchecked) and 'MSDP_st_server' (checked).

Storage unit properties on target

This screenshot shows the 'New Storage Unit' dialog for the target storage unit, 'STU_CCat'. The 'Storage unit type' is 'Disk'. The 'Disk type' is 'PureDisk'. Under 'Properties and Server Selection', the 'Select disk pool' is 'DP_MSDP_CCat'. The 'Media server' section has 'Only use the following media servers' selected. The 'Media Servers' list shows 'CCat_st_server' (checked) and 'MSDP_st_server' (unchecked).

Note: Do not uncheck the box for the Cloud Catalyst host credentials. Only uncheck the box when directed by technical support staff.

Example: In [Figure 6-4](#), do not uncheck the box next to **CCat_st_server** when configuring the storage unit properties on target.

See [Note about credentials on media servers](#)


Example target controlled configuration: Source media storage server is not RHEL or is pre-8.1

In [Figure 6-5](#), the target media server is configured for a target control duplication. Note that the source is Windows and has NetBackup 8.0 installed.

Two configuration items allow for target controlled duplication:

- The source (the MSDP storage server) allows access from the target media server. This is configured in the storage server properties of each.
- The storage unit settings on the source and the target specifically point to the target media server. (Enable **Only use the following media servers** and select the target media server.)

Figure 6-5 Target controlled configuration



Note: Do not uncheck the box for the Cloud Catalyst host credentials. Only uncheck the box when directed by technical support staff.

Example: In [Figure 6-5](#), do not uncheck the box next to **CCat_st_server** when configuring the storage unit properties on target.

See [Note about credentials on media servers](#)

Configuring a Cloud Catalyst storage server as the source for optimized duplication

Starting from NetBackup 8.2, a Cloud Catalyst storage server is supported as the source for optimized duplication.

See [“Configuring a Cloud Catalyst storage server as the target for the deduplications from MSDP storage servers”](#) on page 267.

When you configure a storage lifecycle policy for the backup operation, you can select the storage unit that was created for the Cloud Catalyst storage server. The following table lists the version requirements for configuring source control or target control optimized duplication for Cloud Catalyst.

Table 6-8 Version requirements for optimized duplication using Cloud Catalyst

Optimized duplication source	Optimized duplication target	Requirements for source control		Requirements for target control	
		Source	Target	Source	Target
Media Server Deduplication Pool	Cloud Catalyst storage server	RHEL 7.3 or later and NetBackup 8.1 or later.	RHEL 7.3 or later and NetBackup 8.1 or later.	None	RHEL 7.3 or later and NetBackup 8.1 or later.
Cloud Catalyst storage server	Media Server Deduplication Pool	RHEL 7.3 or later and NetBackup 8.2 or later.	None	RHEL 7.3 or later and NetBackup 8.2 or later.	RHEL 7.3 or later and NetBackup 8.2 or later.
Cloud Catalyst storage server	Cloud Catalyst storage server	RHEL 7.3 or later and NetBackup 8.1 or later.	RHEL 7.3 or later and NetBackup 8.1 or later.	RHEL 7.3 or later and NetBackup 8.1 or later.	RHEL 7.3 or later and NetBackup 8.1 or later.

Decommissioning Cloud Catalyst cloud storage

Perform the following procedure before manually deleting any data from the cloud.

To decommission Cloud Catalyst cloud storage

- 1 Delete the storage unit that uses the disk pool.
- 2 Delete the disk pool.

See [“Deleting a Media Server Deduplication Pool”](#) on page 361.

- 3 Delete the Cloud Catalyst storage server.
- 4 Run the `esfs_cleanup.sh` script on the media server to remove all Cloud Catalyst-related information from the server.

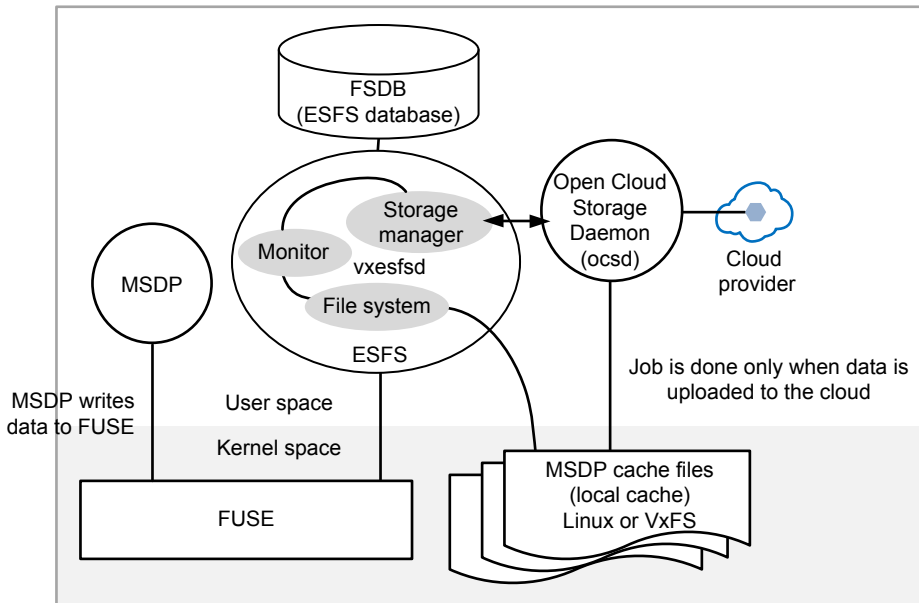
```
/usr/opensv/esfs/scripts/esfs_cleanup.sh
```

- 5 Any remaining data on the cloud storage may now be manually removed.

NetBackup Cloud Catalyst workflow processes

Figure 6-6 shows the workflow on a Cloud Catalyst storage server for uploading data to the cloud.

Figure 6-6 Workflow on a Cloud Catalyst storage server



Note that a single Cloud Catalyst storage server can only write to one cloud provider and to one bucket within that cloud provider.

Filesystem in Userspace (FUSE)

The Filesystem in Userspace (FUSE) forwards data from MSDP to the File System Database (FSDb).

File System Database (FSDB)

The File System Database (FSDB) tracks and stores metadata about all of the files that are written to the NetBackup Extendable Storage File System (ESFS).

NetBackup Extendable Storage File System (ESFS)

NetBackup Cloud Catalyst uses the NetBackup Extendable Storage File System Service (`vxesfsd`) and its subcomponents to move and manage files in the local cache directory and the cloud.

ESFS uses daemons to perform the following functions in its database:

- Veritas NetBackup Extendable Storage File System Service daemon (`vxesfsd`): This process is the primary file system daemon. This process is responsible for writing data into the MSDP cache files.
- Open Cloud Storage Daemon (`ocsd`): This daemon is responsible for interacting with the cloud. Under normal circumstances, there is only one `ocsd` process running.

The `ocsd` daemon produces the following log:

```
/usr/openv/netbackup/logs/esfs_storage
```

`vxesfsd` includes three subcomponents that perform various tasks as part of ESFS:

- File system component: Interacts with the ESFS database.
- Monitor component: The monitor component checks the file queue for files to be uploaded and checks the local cache directory for cache eviction purposes, among other activities.
- Storage manager component: Notifies `ocsd` processes when there's work to be done. Some of the activities that this component manages are synchronous (downloads from the cloud), asynchronous (uploads to and deletions in the cloud), and shared memory.

The `vxesfsd` daemon produces the following logs in `/usr/openv/netbackup/logs/`:

<code>esfs_filesystem</code>	Logs the FUSE operations.
<code>esfs_storagemanager</code>	Logs the interactions of the ESFS storage manager with <code>vxesp</code> .
<code>esfs_monitor</code>	Logs the job status updates and periodic tasks such as cache eviction and FSDB backups.
<code>esfs_database</code>	Logs the FSDB actions that are specific to the ESFS database.

MSDP cache files

As part of the Cloud Storage Server Configuration Wizard, the administrator configures a local cache directory (`local_cache_dir`). The directory must be on either a Linux file system or, in the case of an appliance, a VxFS file system.

The directory is then split into two:

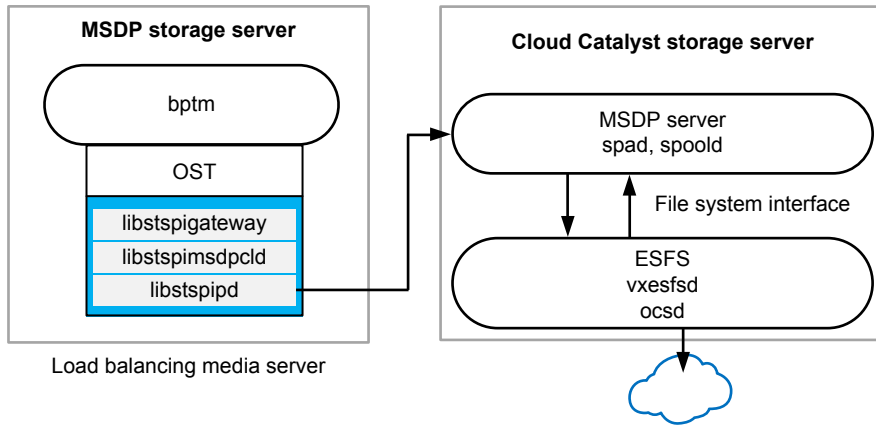
- A cache directory: `local_cache_dir/cache/userdata`
ESFS writes out the data and the metadata in this location. This directory is mapped directly to the storage directory (`local_cache_dir/storage`). Two copies of each file do not exist on this media server.
- A storage directory: `local_cache_dir/storage`
This path is the mount point to the ESFS file system. This path is what MSDP knows as its storage directory.

NetBackup periodically performs data eviction tasks in the cache directory to create space for new backups. If necessary, the administrator can change when or how often data eviction occurs. These options are configured in the `esfs.json` file.

See [“About the Cloud Catalyst esfs.json configuration file”](#) on page 271.


Backup workflow

Figure 6-7 Backup data workflow for a Cloud Catalyst storage server



Optimized duplication workflow

Figure 6-8 Optimized duplication workflow with target as a Cloud Catalyst storage server



Disaster recovery for Cloud Catalyst

NetBackup Cloud Catalyst uses a media server that is configured as a Cloud Catalyst storage server for deduplication to the cloud. The media server can be either a Cloud Catalyst appliance or a non-appliance host. The following link is a document that contains disaster recovery procedures in environments where NetBackup Cloud Catalyst is used.

The document contains two disaster recovery procedures and they have the following conditions:

- The procedures assume that an MSDP catalog backup exists. See [“About protecting the MSDP catalog”](#) on page 233.
- The procedures recover the MSDP cloud storage server, whether it was an appliance or a non-appliance host.

If you encounter any issues with disaster recovery, we recommend that you navigate to support.veritas.com. The Veritas Support site offers additional information to help you complete the process.

https://www.veritas.com/content/support/en_US/doc/Cloud_Catalyst_dr_82

About image sharing in cloud using Cloud Catalyst

Image sharing, which was earlier known as Automated disaster recovery (DR), provides a self-describing storage solution over Cloud Catalyst. Cloud Catalyst with Image sharing in cloud is a self-describing storage server. Cloud Catalyst without Image sharing in cloud is not a self-describing storage server.

Image sharing helps users with an easy and a visualized way to manage and provision images in cloud object storage and even the ability to convert backed up VMs as AWS instances in certain scenarios.

This topic contains the following sections:

- [Important features of image sharing](#)
- [Things to consider before you use image sharing](#)
- [Configure image sharing](#)
- [Using image sharing using NetBackup Web UI](#)
- [Using image sharing with the `nbimageshare` command](#)
- [Manual KMS key transfer in Image sharing in case of NetBackup KMS](#)
- [Manual steps in image sharing in case of external KMS](#)

- [Additional information about image sharing](#)

Important features of image sharing

- In a situation where Cloud Catalyst backed up the deduplicated data to cloud, but the NetBackup catalog was available only on the on-premises NetBackup server. There, the data cannot be restored from the cloud without the on-premises NetBackup server.
Image sharing in cloud uploads the NetBackup catalog along with the backup images and lets you restore data from the cloud without the on-premises NetBackup server.
- You can launch an all-in-one NetBackup in the cloud on demand called the cloud recovery host, and recover the backup images from cloud.
- Image sharing discovers the backup images that are stored in AWS S3 through the REST APIs, recovers the NetBackup catalog, and restores the images.
- You can use command line options or NetBackup Web UI that have the function as REST APIs.

Things to consider before you use image sharing

- Before you install NetBackup, create an instance based on RHEL 7.3 or later (up to RHEL 8.0) in AWS. You can also set up a computer based on RHEL 7.3 or later (up to RHEL 8.0). The recommendation is that the instance has more than 64 GB of memory, 8 CPUs.
- The HTTPS port 443 is enabled.
- Change host name to the server's FQDN.
- Add the following items in the `/etc/hosts` file:
"External IP" "Server's FQDN"
"Internal IP" "Server's FQDN"
For a computer, add the following items in the `/etc/hosts` file:
"IP address" "Server's FQDN"
- For an instance in AWS, change the search domain order in the `/etc/resolv.conf` file to search external domains before internal domains.
- NetBackup should be an all-in-one setup.
Refer to the *NetBackup Installation Guide* for more information.

Configure image sharing

After installing NetBackup, you can run the `ims_system_config.py` script to configure image sharing.

The path to access the command is: `/usr/openv/pdde/pdag/scripts/.`

Use the following command to run the `ims_system_config.py` script:

Amazon Web Service cloud provider:

```
ims_system_config.py -k <AWS_access_key> -s <AWS_secret_access_key>
-b <name_S3_bucket>
```

If you have configured IAM role in the EC2 instance, use the following command:

```
-python /usr/openv/pdde/pdag/scripts/ims_system_config.py
-k dummy -s dummy -b <name_S3_bucket>
```

Microsoft Azure cloud provider:

```
ims_system_config.py -cp 2 -k <key_id> -s <secret_key> -b <container_name>
```

Other S3 compatible cloud provider (For example, Hitachi HCP):

If Cloud Instance has been existed in NetBackup, use the following command:

```
ims_system_config.py -cp 3 -t PureDisk -k <key_id> -s <secret_key>
-b <bucket_name> -bs <bucket_sub_name> -c <Cloud_instance_name> [-p
<mount_point>]
```

Or use the following command:

```
ims_system_config.py -cp 3 -t PureDisk -k <key_id> -s <secret_key>
-b <bucket_name> -pt <cloud_provider_type> -sh <s3_hostname> -sp
<s3_http_port> -sps <s3_https_port> -ssl <ssl_usage> [-p
<mount_point>]
```

Example for HCP provider:

```
ims_system_config.py -cp 3 -t PureDisk -k xxx -s xxx -b emma -bs
subtest -pt hitachicp -sh yyy.veritas.com -sp 80 -sps 443 -ssl 0
```

Description: (Specify the following options to use HCP cloud)

-cp 3: Specify third-party S3 cloud provider that is used.

-pt hitachicp: Specify cloud provider type as **hitachicp** (HCP LAN)

-t PureDisk_hitachicp_rawd: Specify storage server type as **PureDisk_hitachicp_rawd**

-sh <s3_hostname>: Specify HCP storage server host name

-sp <s3_http_port>: Specify HCP storage server HTTP port (Default is 80)

-sps <s3_https_port>: Specify HCP storage server HTTP port (Default is 443)

-ssl <ssl_usage>: Specify whether to use SSL. (0- Disable SSL. 1- Enable SSL. Default is 1.) If SSL is disabled, it uses <s3_http_port> to make connection to <s3_hostname>. Otherwise, it uses <s3_https_port>.

Using image sharing using NetBackup Web UI

You can access NetBackup Web UI to use image sharing. For more information, refer to the *Using image sharing from the NetBackup Web UI* topic in the *NetBackup Web UI Administrator's Guide*.

Using image sharing with the `nbimageshare` command

You can use the `nbimageshare` command to use image sharing.

Run the `nbimageshare` command to list and import the virtual machine and standard images and then recover the virtual machines.

The path to access the command is: `/usr/opensv/netbackup/bin/admincmd/`

For more information about the `nbimageshare` command, refer to the *NetBackup Commands Reference Guide*.

The following table lists the steps for image sharing and the command options:

Table 6-9 Steps for image sharing and the command options

Step	Command
Log on to NetBackup	<pre>nbimageshare --login <username> <password></pre> <pre>nbimageshare --login -interact</pre>
List all the backup images that are in the cloud	<pre>nbimageshare --listimage</pre> <p>Note: In the list of images, the increment schedule type might be differential incremental or cumulative incremental.</p>

Table 6-9 Steps for image sharing and the command options (*continued*)

Step	Command
Import the backup images to NetBackup	<p>Import a single image:</p> <pre>nbimageshare --singleimport <client> <policy> <backupID></pre> <p>Import multiple images:</p> <pre>nbimageshare --batchimport <image_list_file_path></pre> <p>Note: The format of the <code>image_list_file_path</code> is same as the output of "list images".</p> <p>The multiple images number must be equal to or less than 64.</p> <p>You can import an already imported image. This action does not affect the NetBackup image catalog.</p>
Recover the VM as an AWS EC2 instance	<pre>nbimageshare --recovervm <client> <policy> <backupID></pre> <ul style="list-style-type: none"> Only VM images are supported. The AWS account must have the following read and write permissions to S3: <pre>"ec2:CreateTags" "ec2:DescribeImportImageTasks" "ec2:ImportImage" "iam:ListRolePolicies" "iam:ListRoles" "iam:GetRole" "iam:GetRolePolicy" "iam:CreateRole" "iam:PutRolePolicy"</pre>

Manual KMS key transfer in Image sharing in case of NetBackup KMS

When KMS encryption is enabled, you can share the images in S3 bucket to the cloud recovery host with manual KMS key transfer.

On-premises side:

- Storage server: Find the key group name for the given Storage server

Find `contentrouter.cfg` in `/etc/pdregistry.cfg`

Find key group name is in `contentrouter.cfg` under `[KMSOptions]`

(Example `KMSKeyGroupName=amazon.com:test1`)

2. NetBackup master server: Exports the key group with a passphrase to a file:

```
/usr/opensv/netbackup/bin/admincmd/nbkmsutil -export -key_groups
<key-group-name> -path <key file path>
```

Cloud recovery host (cloud side):

1. Copy the exported key to the cloud recovery host
2. Config KMS server

```
/usr/opensv/netbackup/bin/nbkms -createemptydb
/usr/opensv/netbackup/bin/nbkms
/usr/opensv/netbackup/bin/nbkmscmd -discovernbkms -autodiscover
```

3. Import keys to KMS service.

```
/usr/opensv/netbackup/bin/admincmd/nbkmsutil -import -path <key
file path> -preserve_kgname
```

4. Configure the cloud recovery host with `ims_system_config.py`

On-premises KMS key changes

In case of KMS key changes for the given group for on-premises storage server after the cloud recovery host is set up, you must export the key file from on-premises KMS server and import that key file on the cloud recovery host.

1. On-premises NetBackup master server: Exports the key group with a passphrase to a file

```
/usr/opensv/netbackup/bin/admincmd/nbkmsutil -export -key_groups
<key-group-name> -path <key file path>
```

2. Cloud recovery host:

```
/usr/opensv/netbackup/bin/admincmd/nbkmsutil -deletekg -kgname
<key-group-name> -force

/usr/opensv/netbackup/bin/admincmd/nbkmsutil -import -path <key
file path> -preserve_kgname
```

Manual steps in image sharing in case of external KMS

If an on-premises storage server is configured to use keys from external KMS server, then make sure that the same KMS server is configured on the cloud

recovery host before running `ims_system_config.py`. To know more about configuring an external KMS server in NetBackup, refer to *NetBackup Security and Encryption Guide*.

Make sure that the external KMS server is reachable from the cloud recovery host on a specific port.

Additional information about image sharing

- After the image is imported to cloud, the image catalog still exists on the cloud. If the image is expired on the on-premises storage, then restoring the image on the cloud fails even though the image catalog exists on the cloud.

- If the image expires on the cloud storage, the image catalog in the cloud is removed but the image data in the bucket is not removed.

- You can restore any image that you import. For the recover the VM as an AWS EC2 instance option, you can recover only the VM images that are full backup images or accelerator-enabled incremental backup images to cloud.

- Image sharing supports many policy types from NetBackup 8.2 or later. In the scenarios, CloudCatalyst, where the images are shared, must have a new installation of NetBackup 8.2 or later.

See the NetBackup compatibility lists for the latest information on the supported policy types.

- After the image sharing is configured, the storage server is in a read-only mode.

- For information on the VM recovery limitations, refer to the AWS VM import information in AWS help.

- You can configure the maximum active jobs when the images are imported to cloud storage.

Modify the file path `/usr/opensv/var/global/wsl/config/web.conf` to add the configuration item as `imageshare.maxActiveJobLimit`.

For example, `imageshare.maxActiveJobLimit=16`.

The default value is 16 and the configurable range is 1 to 100.

If the import request is made and the active job count exceeds the configured limit, the following message is displayed:

"Current active job count exceeded active job count limitation".

- The images that are direct to cloud storage can be shared.
In optimized deduplication or AIR cascading scenarios, only the images in Cloud Catalyst that has optimized deduplication or has an AIR target can be shared. If Cloud Catalyst is not set for optimized deduplication or is not an AIR target, you cannot use image sharing. If Amazon Glacier is enabled in Cloud Catalyst, you cannot use image sharing.

In these scenarios to disable image sharing:

Modify the `<install_directory>/etc/puredisk/spa.cfg` file and add the following configuration item:

```
EnableIMandTIR=false
```

- Regarding the errors about role policy size limitation:
Errors that occur when the role policy size exceeds the maximum size is an AWS limitation. You can find the following error in a failed restore job:

```
"error occurred (LimitExceeded) when calling the PutRolePolicy operation:  
Maximum policy size of 10240 bytes exceeded for role vmimport"
```

Workaround:

- You can change the maximum policy size limit for the `vmimport` role.
- You can list and delete the existing policies using the following commands:

```
aws iam list-role-policies --role-name vmimport  
aws iam delete-role-policy --role-name vmimport --policy-name  
<bucketname> -vmimport
```

- The recover operation includes AWS import process. Therefore, a vmdk image cannot be recovered concurrently in two restore jobs at the same time.
- The image sharing feature can recover the virtual machines that satisfy the Amazon Web Services VM import prerequisites.

For more information about the prerequisites, refer to the following article:

https://docs.aws.amazon.com/vm-import/latest/userguide/vmie_prereqs.html

- If you cannot obtain the administrator password to use an AWS EC2 instance that has a Windows OS, the following error is displayed:

```
Password is not available. This instance was launched from a custom  
AMI, or the default password has changed. A password cannot be  
retrieved for this instance. If you have forgotten your password,  
you can reset it using the Amazon EC2 configuration service. For  
more information, see Passwords for a Windows Server Instance.
```

This error occurs after the instance is launched from an AMI that is converted using image sharing.

For more information, refer to the following articles:

- [Amazon Elastic Compute Cloud Common Messages](#)
- [How to migrate your on-premises domain to AWS Managed Microsoft AD using ADMT](#)
- You cannot cancel an import job on the cloud recovery host.

- If there is data optimization done on the on-premises image, you might not be able to restore the image that you have imported on the cloud recovery host. You can expire this image, import it again on the image-sharing server, and then restore the image.
- After the backup job, duplication job, or AIR import job completes, you can import the images on a cloud recovery host.

MSDP cloud support

This chapter includes the following topics:

- [About MSDP cloud support](#)
- [Creating a cloud storage unit](#)
- [Updating cloud credentials for a cloud LSU](#)
- [Updating encryption configurations for a cloud LSU](#)
- [Deleting a cloud LSU](#)
- [Backup data to cloud by using cloud LSU](#)
- [Duplicate data cloud by using cloud LSU](#)
- [Configuring AIR to use cloud LSU](#)
- [About backward compatibility support](#)
- [About the configuration items in cloud.json, contentrouter.cfg and spa.cfg](#)
- [About the tool updates for cloud support](#)
- [About the disaster recovery for cloud LSU](#)
- [About Image Sharing using MSDP cloud](#)

About MSDP cloud support

In this release, NetBackup MSDP cloud support is enhanced to provide a flexible, scalable, high performing, and easy to configure solution, that enables you to leverage cloud storage more efficiently.

Here are the highlights of this feature:

- One MSDP storage server can be configured to support multiple storage targets, including one local storage target and zero or more cloud storage targets. You can move data to local and multiple cloud targets simultaneously.
- The cloud targets can be from the same or different providers, either public, or private. For example, AWS, Azure, HCP, etc.
- The cloud targets can be added on demand after the MSDP server is configured and active.
- Multiple cloud targets can coexist in a single cloud bucket or multiple buckets that are distributed in a single or different cloud providers.
- The data and metadata for local storage and multiple cloud targets are isolated to support Multiple Tenant usage.
- Optimized deduplication is supported within one MSDP server scope so that data can be stored to local storage first and then duplicated to cloud targets in the same media server.
- Disaster recovery from the cloud targets is enhanced and more straightforward.
- Feature is well integrated with the MSDP cluster solution.

Based on the OpenStorage Technology (OST), the new architecture uses multiple Logical storage unit (LSU) to manage and move data. These LSUs can be customized independently to meet different customer requirements. For example, as pure local target (same as MSDP in NetBackup 8.2 or earlier), or local target plus one or more cloud targets.

Starting with NetBackup 8.3, you can configure MSDP from the NetBackup Web UI. You can refer to the NetBackup Web UI documentation for more details.

This chapter focuses on how to use the command line interface to configure MSDP.

Note: To enable OCSD logging information or MSDP cloud, add `loglevel=3` in the section `[Symantec/PureDisk/OpenCloudStorageDaemon]` in `/etc/pdregistry.cfg` on media server and restart the services.

Check the logs at `/usr/opensv/netbackup/esfs_storage/`.

Operating system requirement for configuration

Cloud LSUs can be configured on the storage servers running on Red Hat Linux Enterprise or CentOS platforms. No platform limitations for clients and load balancing servers.

Limitations

- Instant Access for cloud LSU is not supported.
- Universal share for cloud LSU is not supported.
- Accelerator for cloud LSU of AWS Glacier and Deep Archive is not supported.
- Cloud DR for cloud LSU of AWS Glacier and Deep Archive is not supported if the storage server name changes.
- AIR for cloud LSU of AWS Glacier and Deep Archive is not supported.
- Classic AIR for cloud LSU as replication target is not supported.
- Synthetic backup for cloud LSU of AWS Glacier and Deep Archive is not supported.
- Snowball is not supported.
- FIPS is not supported.
- Multi-threaded Agent must be disabled when a Client-Direct backup is used by NetBackup clients that have NetBackup version earlier than 8.3.
- If you select a load balancing media server that has NetBackup version earlier than 8.3, then the cloud LSUs are not listed. Even if you select cloud LSUs with media servers that have NetBackup version earlier than 8.3, the backups can fail.

Creating a cloud storage unit

Use the NetBackup Web UI or the command line to create a cloud storage unit.

The following steps describe the method to create a cloud storage unit using the command line:

- 1 Create an MSDP storage server.

See [“Configuring MSDP server-side deduplication”](#) on page 55.

- 2 Create a cloud instance alias.

For example:

Example 1: Creating an Amazon S3 cloud instance alias

```
# /usr/opensv/netbackup/bin/admincmd/csconfig cldinstance -as -in  
amazon.com -sts <storage server> -lsu_name <lsu name>
```

Example 2: Creating an Amazon Glacier cloud instance alias

```
# /usr/opensv/netbackup/bin/admincmd/csconfig cldinstance -as -in  
amazon.com -sts <storage server> -lsu_name <lsu name>  
-storage_class GLACIER
```

The cloud alias name is `<storage server>_<lsu name>`, and is used to create a bucket.

- 3 Create a new bucket (Optional)

For example:

```
# /usr/opensv/netbackup/bin/nbclutil -createbucket -storage_server  
<storage server>_<lsu name> -username <cloud user> -bucket_name  
<bucket name>
```

4 Create a configuration file, then run `nbdevconfig` command.

Configuration file content for adding a new cloud LSU:

Configuration setting	Description
V7.5 "operation" "add-lsu-cloud" string	Specifies the value "add-lsu-cloud" for adding a new cloud LSU.
V7.5 "lsuName" " " string	Specifies the LSU name.
V7.5 "lsuCloudUser" " " string	Specifies the cloud user name.
V7.5 "lsuCloudPassword" " " string	Specifies the cloud password.
V7.5 "lsuCloudBucketName" " " string	Specifies the cloud bucket name.
V7.5 "lsuCloudBucketSubName" " " string	Multiple cloud LSUs can use the same cloud bucket, this value distinguishes different cloud LSUs.
V7.5 "lsuEncryption" " " string	Optional value, default is NO. Sets the encryption property for current LSU.
V7.5 "lsuKmsEnable" " " string	Optional value, default is NO. Enables KMS for current LSU.
V7.5 "lsuKmsKeyGroupName" " " string	Optional value. Key group name is shared among all LSUs. Key group name must have valid characters: A-Z, a-z, 0-9, _ (underscore), - (hyphen), : (colon), . (period), and space.
V7.5 "lsuKmsServerName" " " string	Optional value. KMS server name is shared among all LSUs.
V7.5 "lsuKmsServerType" " " string	Optional value.

Example 1: Configuration file with encryption disabled

```
V7.5 "operation" "add-lsu-cloud" string
V7.5 "lsuName" "s3amazon1" string
V7.5 "lsuCloudUser" "CUCU" string
V7.5 "lsuCloudPassword" "cpcp" string
V7.5 "lsuCloudBucketName" "bucket1" string
V7.5 "lsuCloudBucketSubName" "sub1" string
```

Example 2: Configuration file with encryption enabled

```
V7.5 "operation" "add-lsu-cloud" string
V7.5 "lsuName" "s3amazon2" string
V7.5 "lsuCloudUser" "CUCU" string
V7.5 "lsuCloudPassword" "cpcp" string
V7.5 "lsuCloudBucketName" "bucket1" string
V7.5 "lsuCloudBucketSubName" "sub2" string
V7.5 "lsuEncryption" "YES" string
V7.5 "lsuKmsEnable" "YES" string
V7.5 "lsuKmsKeyGroupName" "test" string
V7.5 "lsuKmsServerName" "test" string
```

Note: All encrypted LSUs in one storage server must use the same `keygroupname` and `kmsservername`. If you use the `nbdevconfig` command to add a new encrypted cloud Logical storage unit (LSU) and an encrypted LSU exists in this MSDP, the `keygroupname` must be the same as the `keygroupname` in the previous encrypted LSU.

For more information, See [“About MSDP Encryption using NetBackup KMS service”](#) on page 76.

Create a configuration file and then run the following `nbdevconfig` command:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server <storage server> -stype PureDisk -configlist
<configuration file path>
```

Note: The parameter `<storage server>` must be the same as the parameter `<storage server>` in Step 2.

5 Create disk pool by using the `nbdevconfig` command.

For example:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -previewdv
-storage_servers <storage server name> -stype PureDisk | grep
<LSU name> > /tmp/dvlist

# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -createdp -dp
<disk pool name> -stype PureDisk -dvlist /tmp/dvlist
-storage_servers <storage server name>
```

Note: You can also create the disk pool can from the NetBackup Web UI or NetBackup Administration Console.

6 Create storage unit by using `bpstuadd` command.

For example:

```
# /usr/opensv/netbackup/bin/admincmd/bpstuuadd -label <storage unit
name> -odo 0 -dt 6 -dp <disk pool name> -nodevhost
```

Note: You can also create the storage unit from the NetBackup Web UI or NetBackup Administration Console.

Updating cloud credentials for a cloud LSU

To update cloud credentials for a cloud LSU, you can create a configuration file and then run `nbdevconfig` command.

Configuration file contents for updating cloud credential are as follows:

Configuration setting	Description
V7.5 "operation" "update-lsu-cloud" string	Use the value "update-lsu-cloud" to update some cloud LSU parameters.
V7.5 "lsuName" " " string	Specifies the LSU name.
V7.5 "lsuCloudUser" " " string	Specifies the cloud user name.
V7.5 "lsuCloudPassword" " " string	Specifies the cloud password.

For example:

```
V7.5 "operation" "update-lsu-cloud" string
V7.5 "lsuName" "s3amazon1" string
V7.5 "lsuCloudAlias" "amazon_1" string
V7.5 "lsuCloudUser" "ChangedCloudUser" string
V7.5 "lsuCloudPassword" "changedpassword" string
```

After creating the configuration file, run the following command:

```
# /usr/openv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server <storage server> -stype PureDisk -configlist
<configuration file path>
```

Updating encryption configurations for a cloud LSU

To enable KMS encryption configurations for a Cloud LSU, you can create a configuration file and then run `nbdevconfig` command.

Configuration file contents for updating encryption configurations are as follows:

Configuration setting	Description
V7.5 "operation" "update-lsu-cloud" string	You can only update the KMS status from disabled to enabled.
V7.5 "lsuName" " " string	Specifies the LSU name.
V7.5 "lsuKmsEnable" "YES" string	Specifies the KMS status for the cloud LSU.
V7.5 "lsuKmsServerName" "" string	Optional value. KMS server name that is shared among all LSUs.
V7.5 "lsuKmsKeyGroupName" "" string	Optional value. Key group name that is shared among all LSUs. Key group name must have valid characters: A-Z, a-z, 0-9, _ (underscore), - (hyphen), : (colon), . (period), and space.

Example to enable KMS status from disabled status to enabled status for cloud LSU "s3amazon":

```
V7.5 "operation" "update-lsu-cloud" string
V7.5 "lsuName" "s3amazon" string

V7.5 "lsuKmsEnable" "YES" string
```

```
V7.5 "lsuKmsServerName" "XXX" string
V7.5 "lsuKmsKeyGroupName" "XXX" string
```

Note: All encrypted LSUs in one storage server must use the same `keygroupname` and `kmsservername`. If you use the `nbdevconfig` command to add a new encrypted cloud Logical storage unit (LSU) and an encrypted LSU exists in this MSDP, the `keygroupname` must be the same as the `keygroupname` in the previous encrypted LSU.

For more information, See [“About MSDP Encryption using NetBackup KMS service”](#) on page 76.

After creating the configuration file, run the following command:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server <storage_server> -stype PureDisk -configlist
<configuration_file_path>
```

Deleting a cloud LSU

Use the following steps carefully to remove an MSDP cloud LSU:

- 1 Expire all images of the cloud LSU in NetBackup.
- 2 Remove the storage unit and disk pool of this MSDP cloud LSU.
- 3 To delete a cloud LSU, `storageId` and `CachePath` are needed.

Run following command to get the information of one cloud LSU:

```
/usr/opensv/pdde/pdcr/bin/pddecfg -a listcloudlsu
dsid, lsuname, storageId, CachePath
3, S3Volume, server1_ S3Volume/cloud-bucket1/sub1, /msdp/data/ds_3
4, S3Volume2, server1_ S3Volume2/cloud-bucket1/sub2,
/msdp/data/ds_4
```

Here the `storageId` of the cloud LSU is “server1_ S3Volume/cloud-bucket1/sub1” and `CachePath` of the cloud LSU is “/msdp/data/ds_3”

- 4 Run `CRQP` to make sure no `tlog` entries are present in `<msdp_storage_path>/spool` folder and `<msdp_storage_path>/queue` folder.

- 5 Delete LSU configurations in `spad` by using `nbdevconfig` command.

Configuration file contents for deleting an MSDP cloud LSU configuration are as follows:

Configuration setting	Description
V7.5 "operation" "delete-lsu-cloud" string	The value "delete-lsu-cloud" for deleting the MSDP cloud LSU configurations in <code>spad</code> .
V7.5 "lsuName" " " string	Specifies the LSU name.

For example:

```
V7.5 "operation" "delete-lsu-cloud" string
V7.5 "lsuName" "s3amazon1" string
```

After creating the configuration file, run the following command:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server <storage_server> -stype PureDisk -configlist
<configuration_file_path>
```

- 6 Stop the MSDP service and its monitor service.

```
# /usr/opensv/netbackup/bin/nbsvcmon -terminate
# /usr/opensv/pdde/pdconfigure/pdde stop
```

- 7 Delete LSU configurations in `spoold` using the following command:

```
# /usr/opensv/pdde/pdcr/bin/spoold --removepartition <storageId>
```

- 8 Remove the cache folder by using the following command (Optional):

```
# rm -r <CachePath>
```

- 9 Remove the entire sub-bucket folder in cloud. (Optional)

- 10 Start the MSDP service and its monitor service.

```
# /usr/opensv/pdde/pdconfigure/pdde start
# /usr/opensv/netbackup/bin/nbsvcmon
```

Backup data to cloud by using cloud LSU

Run the following steps to backup data to cloud LSU:

- Create a cloud LSU and related disk pool and storage unit (cloud storage unit).
- Create a backup policy and use cloud storage unit as policy storage.

- Run the backup and the data is written to cloud storage.

You can create and back up to multiple Cloud LSUs at the same storage server.

Duplicate data cloud by using cloud LSU

Run the following steps to duplicate backup images from local MSDP to cloud LSU:

- Configure an MSDP storage server and create a disk pool by using “PureDiskVolume” and then create a storage unit (local storage unit).
- Create a cloud LSU and related disk pool and storage unit (cloud storage unit).
- Create a Storage Lifecycle Policy and add “Backup” and “Duplication” values. Data is backs up to local storage unit and then duplicates to a cloud storage unit.
- Create a backup policy and use storage lifecycle policy as policy storage.
- Run the backup and the data is written to local storage and then duplicated to cloud storage.

Duplication can also be done from cloud LSU to local MSDP and between two cloud LSUs.

Configuring AIR to use cloud LSU

The following steps describe the tasks that are required to replicate the backup images from one LSU to another in a different NetBackup domain:

- See [“Configuring MSDP replication to a different NetBackup domain”](#) on page 158.
- Configure a trust relationship with the target NetBackup domain.
See [“About trusted master servers for Auto Image Replication”](#) on page 167.
- Add an LSU in the remote storage server as a replication target.
To add a replication target in a different NetBackup domain, you can use NetBackup Web UI or use the command line interface.

1 Create a configuration file for adding a replication target.

Configuration file content for adding a replication target:

Configuration setting	Description
V7.5 "operation" " " string	The value must be "set-replication" for adding a new replication target.
V7.5 "rephostname" " " string	Specifies the replication target's host name.
V7.5 "relogin" " " string	Specifies the replication target storage server's user name.
V7.5 "repasswd" " " string	Specifies the replication target storage server's password.
V7.5 "repsourcevolume" " " string	Specifies the replication source volume name.
V7.5 "reptargetvolume" " " string	Specifies the replication target volume name.

Example:

```
[root@sourceserver~]# cat add-replication-local2cloud.txt
V7.5 "operation" "set-replication" string
V7.5 "rephostname" "targetserver1.example.com" string
V7.5 "relogin" "root" string
V7.5 "repasswd" "root" string
V7.5 "repsourcevolume" "PureDiskVolume" string
V7.5 "reptargetvolume" "s3amazn1" string
```

After creating the configuration file, run the `nbdevconfig` command.

For example:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server <storage_server> -stype PureDisk -configlist
<configuration_file_path>
```

2 Run `nbdevconfig` to update the disk volume.

For example:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -updatedv -stype
PureDisk -dp diskpool1 -media_server sourceserver.example.com
```

■ Configure a storage lifecycle policy.

You must create an import SLP in the target domain before you configure an SLP in the source domain.

See [“About storage lifecycle policies”](#) on page 192.

See [“About the storage lifecycle policies required for Auto Image Replication”](#) on page 192.

See [“Creating a storage lifecycle policy”](#) on page 194.

Removing a replication target

Complete the following steps to delete a replication target:

1. Create a configuration file for deleting a replication target.

Configuration file content for deleting a replication target:

Configurtion setting	Description
V7.5 "operation" " " string	The value must be “delete-replication” for deleting a new replication target.
V7.5 "rephostname" " " string	Specifies the replication target's host name.
V7.5 "repsourcevolume" " " string	Specifies the replication source volume name.
V7.5 "reptargetvolume" " " string	Specifies the replication target volume name.

For example:

```
[root@sourceserver~]# cat delete-replication-local2cloud.txt
V7.5 "operation" "delete-replication" string
V7.5 "rephostname" "targetserver1.example.com" string
V7.5 "repsourcevolume" "PureDiskVolume" string
V7.5 "reptargetvolume" "s3amamzon1" string
```

After creating the configuration file, run the `nbdevconfig` command.

For example:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server <storage server> -stype PureDisk -configlist
<configuration file path>
```

2. Run `nbdevconfig` to update the disk volume.

For example:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -updatedv -stype
PureDisk -dp diskpool1 -media_server sourceserver.example.com
```

About backward compatibility support

If an Auto Image Replication (A.I.R) source storage server has NetBackup 8.2 or earlier, the target disk volume cannot be a cloud LSU.

To replicate image from an earlier version of an MSDP server to the cloud LSU of target MSDP server, you need a user name with the cloud LSU name when you add the A.I.R target. The format of the user name with the target cloud LSU is:

```
<username>?LSU=<target cloud LSU>
```

For example, there is a target storage server and user name of the server is `userA` and there is a cloud LSU `s3cloud1` in the target storage server. To replicate image from an old storage server to the cloud LSU of the target server, you can use the following user name while adding the A.I.R target:

```
userA?LSU=s3cloud1
```

You must also create an import SLP to local volume of target storage server in the target master server. Then select the imported SLP while creating the target A.I.R SLP on the source side. When A.I.R runs, the import job in target side shows the policy name as **SLP_No_Target_SLP** in the Activity Monitor, but the data is sent to cloud.

If the NetBackup client version is 8.2 or earlier, the client direct backup from the old client to cloud LSU of one storage server might fail. During the backup if `mtstrmd` is used on the client side, the job fails with a media write error. To disable `mtstrmd` at the client side, open the configuration file `pd.conf` on the client and change the following:

```
MTSTRM_BACKUP_ENABLED = 1 to MTSTRM_BACKUP_ENABLED = 0.
```

The `pd.conf` file is located in the following directories:

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

When doing a client direct backup with a cloud LSU and an old client, the client does only client-side deduplication.

To use cloud LSU, the load balance server of the storage server must not be an earlier version (NetBackup 8.2 or earlier). If there are new and old load balancers, new load balance server is selected automatically to make sure that the job can be done successfully. When you are restoring a backup image on cloud LSU and you select the media server explicitly, the media server selected must not be an earlier version of NetBackup.

About the configuration items in cloud.json, contentrouter.cfg and spa.cfg

The cloud.json file is available at: <STORAGE>/etc/puredisk/cloud.json.

The file has the following parameters:

Parameter	Details	Default value
UseMemForUpload	<p>If it is set to true, the upload cache directory is mounted in memory as tmpfs. It is especially useful for high speed cloud that disk speed is bottleneck. It can also reduce the disk competition with local LSU. The value is set to true if the system memory is enough.</p> <p>The default value is true if there is enough memory available.</p>	true
CachePath	<p>The path of the cache. It is created under an MSDP volume according to the space usage of MSDP volumes. It will reserve some space that local LSU cannot write beyond. Usually you do not need to change this path, unless in some case that some volumes are much freer than others, multiple cloud LSUs may be distributed to the same disk volume. For performance consideration, you may need to change this option to make them distributed to different volumes. This path can be changed to reside in a non-MSDP volume.</p>	NA
UploadCacheGB	<p>It is the maximum space usage of upload cache. Upload cache is a subdirectory named "upload" under CachePath. For performance consideration, it should be set to larger than:</p> <p>$(\text{max concurrent write stream number}) * \text{MaxFileSizeMB} * 2.$</p> <p>So, for 100 concurrent streams, about 13 GB is enough.</p> <p>Note: The initial value of UploadCacheGB in the cloud.json file is the value of CloudUploadCacheSize in the contentrouter.cfg file.</p> <p>When you add a new cloud LSU, the value of UploadCacheGB is equal to CloudUploadCacheSize. You can later change this value in the cloud.json file.</p>	12

Parameter	Details	Default value
DownloadDataCacheGB	<p>It is the maximum space usage of data file, mainly the <code>SO BIN</code> file. The larger this cache, the more data files can reside in the cache. Then there is no need to download these files from cloud when doing restore.</p> <p>Note: The initial value of <code>DataCacheGB</code> in the <code>cloud.json</code> file is the value of <code>CloudDataCacheSize</code> in the <code>contentrouter.cfg</code> file.</p> <p>When you add a new cloud LSU, the value of <code>DataCacheGB</code> is equal to <code>CloudDataCacheSize</code>. You can later change this value in the <code>cloud.json</code> file.</p>	500
DownloadMetaCacheGB	<p>It is the maximum space usage of metadata file, mainly the <code>DO</code> file and <code>SO BHD</code> file. The larger this cache, the more meta files can reside in the cache. Then there is no need to download these files from cloud when doing restore.</p> <p>Note: The initial value of <code>MetaCacheGB</code> in the <code>cloud.json</code> file is the value of <code>CloudMetaCacheSize</code> in the <code>contentrouter.cfg</code> file.</p> <p>When you add a new cloud LSU, the value of <code>MetaCacheGB</code> is equal to <code>CloudMetaCacheSize</code>. You can later change this value in the <code>cloud.json</code> file.</p>	500
MapCacheGB	<p>It is the max space usage of <code>map</code> file that is used for compatibility of MD5 type fingerprint. The larger this cache, the more <code>map</code> files can reside in the cache.</p> <p>Note: The initial value of <code>MapCacheGB</code> in the <code>cloud.json</code> file is the value of <code>CloudMapCacheSize</code> in the <code>contentrouter.cfg</code> file.</p> <p>When you add a new cloud LSU, the value of <code>MapCacheGB</code> is equal to <code>CloudMapCacheSize</code>. You can later change this value in the <code>cloud.json</code> file.</p>	5
UploadConnNum	Maximum number of concurrent connections to the cloud provider for uploading. Increasing this value is helpful especially for high latency network.	60
DataDownloadConnNum	Maximum number of concurrent connections to the cloud provider for downloading data. Increasing this value is helpful especially for high latency network.	40
MetaDownloadConnNum	Maximum number of concurrent connections to the cloud provider for downloading metadata. Increasing this value is helpful especially for high latency network.	40
MapConnNum	Maximum number of concurrent connections to the cloud provider for downloading map.	40

Parameter	Details	Default value
DeleteConnNum	Maximum number of concurrent connections to the cloud provider for deleting. Increasing this value is helpful especially for high latency network.	100
KeepData	Keep uploaded data to data cache. The value always false if UseMem is true.	false
KeepMeta	Keep uploaded meta to meta cache, always false if UseMem is true.	false
ReadOnly	LSU is read only, cannot write and delete on this LSU.	false
MaxFileSizeMB	Max size of bin file in MB.	64
WriteThreadNum	The number of threads for writing data to the data container in parallel that can improve the performance of IO.	2
RebaseThresholdMB	<p>Rebasing threshold (MB), when image data in container less than the threshold, all of the image data in this container will not be used for deduplication to achieve good locality. Allowed values: 0 to half of MaxFileSizeMB, 0 = disabled</p> <p>Adding a new cloud LSU fails if no partition has free space more than the following:</p> $\text{CloudDataCacheSize} + \text{CloudMapCacheSize} + \text{CloudMetaCacheSize} + \text{CloudUploadCacheSize} + \text{WarningSpaceThreshold} * \text{partition size}$ <p>Use the <code>crcontrol --dsstat 2</code> command to check the space of each of the partition.</p>	4
UsableMemoryLimit	<p>Specify the maximum usable memory size in percentage.</p> $\text{MaxCacheSize} + \text{MaxCloudCacheSize} + \text{Cloud in-memory upload cache size}$ <p>must be less than or equal to the value of UsableMemoryLimit</p>	80%

Note: Each Cloud LSU has a cache directory. The directory is created under an MSDP volume that is selected according to the disk space usage of all the MSDP volumes. Cloud LSU reserves some disk space for cache from that volume, and the local LSU cannot utilize more disk space.

The initial reserved disk space for each of the cloud LSU is the sum of values of UploadCacheGB, DownloadDataCacheGB, DownloadMetaCacheGB, and MapCacheGB in the <STORAGE>/etc/puredisk/cloud.json file. The disk space decreases when the caches are used.

There is a Cache options in crcontrol --dsstat 2 output:

```
# crcontrol --dsstat 2

===== Mount point 2 =====

Path = /msdp/data/dpl/lpdvol

Data storage

Raw Size Used Avail Cache Use%

48.8T 46.8T 861.4G 46.0T 143.5G 2%

Number of containers : 3609

Average container size : 252685915 bytes (240.98MB)

Space allocated for containers : 911943468161 bytes (849.31GB)

Reserved space : 2156777086976 bytes (1.96TB)

Reserved space percentage : 4.0%
```

The Cache option is the currently reserved disk space by cloud for this volume. The disk space is the sum of the reserved space for all cloud LSUs that have cache directories on this volume. The actually available space for Local LSU on this volume is Avail - Cache.

The contentrouter.cfg file is available at:
<STORAGE>/etc/puredisk/contentrouter.cfg.

The file has the following parameters:

Parameter	Details	Default value
CloudDataCacheSize	Default data cache size when adding Cloud LSU. Decrease this value if enough free space is not available.	500 GiB

Parameter	Details	Default value
CloudMapCacheSize	Default map cache size when adding Cloud LSU. Decrease this value if enough free space is not available.	5 GiB
CloudMetaCacheSize	Default meta cache size when adding Cloud LSU. Decrease this value if enough free space is not available.	500 GiB
CloudUploadCacheSize	Default upload cache size when adding Cloud LSU. The minimum value is 12 GiB.	12 GiB
MaxCloudCacheSize	Specify the maximum cloud cache size in percentage. It is based on total system memory, swap space excluded.	20 %
CloudBits	The number of top-level entries in the cloud cache. This number is $(2^{\text{CloudBits}})$. Increasing this value improves cache performance, at the expense of extra memory usage. Minimum value = 16, maximum value = 48.	Auto-sized according to <code>MaxCloudCacheSize</code>
DCSCANDownloadTmpPath	While using the <code>dcscan</code> to check cloud LSU, data gets downloaded to this folder. For details, see the <code>dcscan</code> tool in cloud support section. The <code>spa.cfg</code> file is available at: <code><STORAGE>/etc/puredisk/spa.cfg</code> . The file has the following parameters:	disabled

Parameter	Details	Default value
CloudLSUCheckInterval	The check cloud LSU status interval in seconds.	1800

About the tool updates for cloud support

DCSCAN:

`Dcscan` downloads data container from the cloud. The default download path is `<STORAGE>/tmp/DSID_#dsid`, where `#dsid` is dependent on the cloud LSU DSID value. Different cloud storage providers have different DSID values. You do not need to know the DSID value, `dcscan` gets the DSID value automatically. The default download path can be modified in the `contentrouter.cfg` file using the `DCSCANDownloadTmpPath` field.

While using the `dcscan` tool to look at cloud data, `-a` option is disabled, because it downloads all data containers from cloud, it is an expensive operation. The `-fixdo` option is disabled as well, because `dcscan` only downloads data container from the cloud. Others operations are same as the local LSU.

SEEDUTIL:

Seedutil can be used for seeding a backup for a better deduplication rate. It creates links in the <destination client name> directory to all the backup files found in the path <client name>/<policy name> that have <backup ID> in their names. The user needs to know which DSID value the cloud LSU has used. That DSID value needs to be given to the seedutil, to let seedutil know which cloud LSU will seed a client. If you do a seeding for a local LSU, the default DSID is 2, you do not need to give the DSID value. Seedutil cannot seed across different DSIDs.

For example, /usr/opensv/pdde/pdag/bin/seedutil -seed -sclient <source_client_name> -spolicy <source_policy_name> -dclient <destination_client_name> -dsid <dsid_value>.

CRCONTROL:

Using crcontrol -cloudsstat option to show cloud LSU datastore usage. DSID value needs to be given. As cloud storage has unlimited space, the size is hard-coded to 8 PB.

For example:

```
# /usr/opensv/pdde/pdcr/bin/crcontrol --cloudsstat <dsid_value>
***** Data Store statistics *****
Data storage      Raw      Size   Used   Avail  Use%
8.0P              8.0P   80.9G  8.0P   0%
Number of containers      : 3275
Average container size    : 26524635 bytes (25.30MB)
Space allocated for containers : 86868179808 bytes (80.90GB)
Reserved space           : 0 bytes (0.00B)
Reserved space percentage : 0.0%
```

CRSTATS:

Using crstats -cloud -dsid option to show the cloud LSU statistics. DSID value needs to be given. As cloud storage has unlimited space, the size is hard-coded to 8 PB.

For example:

```
#/usr/opensv/pdde/pdcr/bin/crstats --cloud-dsid <dsid_value>
Storage Pool Raw Size=9007199254740992Bytes
Storage Pool Reserved Space=0Bytes
Storage Pool Required Space=0Bytes
Storage Pool Size=9007199254740992Bytes
Storage Pool Used Space=86868179808Bytes
Storage Pool Available Space=9007112386561184Bytes
```

```
Catalog Logical Size=402826059439Bytes
Catalog files Count=3726
Space Allocated For Containers=86868179808Bytes
Deduplication Ratio=4.6
```

PDDECFG:

Using `pddecfg` to list all the cloud LSUs.

For example:

```
/usr/openv/pdde/pdcr/bin/pddecfg -a listcloudlsu
dsid, lsuname, storageId, CachePath
3, S3Volume, amazon_1/cloud-bucket1/sub1, /msdp/data/ds_3
4, S3Volume2, amazon_1/cloud-bucket1/sub2, /msdp/data/ds_4
```

About the disaster recovery for cloud LSU

If the disk on which the NetBackup software resides or the disk on which the deduplicated data resides fails, you can use the following steps to recover the system and data depending on different scenarios.

After recovery, your NetBackup deduplication environment functions normally. Any valid backup images on the cloud LSU storage are available for restores.

Before you start disaster recovery, ensure that:

- NetBackup media server software is installed and connected with master server successfully.
- The storage server is reinstalled and configured.
- KMS encryption in cloud LSU is enabled and the KMS server is ready.
- Cloud LSU of Amazon S3 Glacier and Deep Archive supports cloud disaster recovery only in Scenario 1 and Scenario 3.

Scenario 1: Storage server name has not changed and the local storage is lost

Step	Task	Procedure
0	Reinstall and configure storage server	See the <i>NetBackup Administrator's Guide, Volume I</i> .
1	Create an empty local LSU	Refer to the "Configure/Reconfigure MSDP local storage" topic.
2	Reuse Cloud LSU	Refer to the "Reuse Cloud LSU" topic.

Scenario 2: Storage server name has changed and the local storage is lost

Step	Task	Procedure
1	Expire the backup images	<p>Expire all backup images that reside on the deduplication disk storage.</p> <p>Warning: Do not delete the images. They are imported back into NetBackup later in this process. If you use the <code>bpexpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>.</p>
2	Delete old storage server-related configurations	<p>Refer to the "Recovering from an MSDP storage server failure" topic.</p> <p>Delete the storage units that use the disk pool.</p> <p>Delete the disk pool.</p> <p>Delete the deduplication storage server.</p> <p>Delete the deduplication host configuration file.</p>
3	Configure new storage server.	<pre>/usr/opensv/netbackup/bin/admincmd/nbdevconfig -creates -storage_server "storage server" -stype PureDisk -media_server "media server" -st 9</pre>
4	Create an empty local LSU.	Refer to the "Configure/Reconfigure MSDP local storage" topic.
5	Reuse cloud LSU.	Refer to the "Reuse Cloud LSU" topic.
6	Create disk pool for cloud LSUs.	<pre>/usr/opensv/netbackup/bin/admincmd/nbdevconfig -createdp -stype PureDisk -dp dpname -storage_server sts_hostname -dvlist filename</pre>
7	Import the images back.	<p>Do two-phase import.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>.</p>

Scenario 3: Storage server name has not changed and the local storage is not lost

Step	Task	Procedure
0	Reinstall and configure storage server	See the <i>NetBackup Administrator's Guide, Volume I</i> .
1	Reuse existing local storage path	Refer to the "Configure/Reconfigure MSDP local storage" topic.
2	Restart storage server.	<pre>/usr/opensv/netbackup/bin/bp.kill_all /usr/opensv/netbackup/bin/bp.start_all</pre>

Scenario 4: Storage server name has changed and local storage is not lost

Step	Task	Procedure
1	Expire the backup images.	<p>Expire all backup images that reside on the deduplication disk storage.</p> <p>Warning: Do not delete the images. They are imported back into NetBackup later in this process. If you use the <code>bpexpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>.</p>
2	Delete old storage server-related configurations.	<p>Refer to the "Recovering from an MSDP storage server failure" topic.</p> <p>Delete the storage units that use the disk pool.</p> <p>Delete the disk pool.</p> <p>Delete the deduplication storage server.</p> <p>Delete the deduplication host configuration file.</p>
3	Configure new storage server.	<pre>/usr/opensv/netbackup/bin/admincmd/nbdevconfig -createsets -storage_server "storage server" -stype PureDisk -media_server "media server" -st 9</pre>
4	Reuse existing local storage path	Refer to the "Configure/Reconfigure MSDP local storage" topic.
5	Restart storage server.	<pre>/usr/opensv/netbackup/bin/bp.kill_all /usr/opensv/netbackup/bin/bp.start_all</pre>
6	Create disk pool for cloud LSUs.	<pre>/usr/opensv/netbackup/bin/admincmd/nbdevconfig -createdp -stype PureDisk -dp dpname -storage_server sts_hostname -dvlist filename</pre>
7	Import the images back.	<p>Do two-phase import.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>.</p>

Common disaster recovery steps

Configure/Reconfigure MSDP local storage

Step	Task	Procedure
1	Delete the deduplication configuration.	<pre>/usr/opensv/pdde/pdconfigure/scripts/installers/PDDE_deleteConfig.sh</pre>

Step	Task	Procedure
2	Delete the NetBackup deduplication Engine credentials on load balancing servers.	<pre>/usr/opensv/volmgr/bin/tpconfig -delete -storage_server sts_hostname -stype PureDisk -sts_user_id root -all_hosts /usr/opensv/volmgr/bin/tpconfig -add -storage_server sts_hostname -stype PureDisk -sts_user_id root -password xxx</pre>
3	Get the config template.	<pre>/usr/opensv/netbackup/bin/admincmd/nbdevconfig -getconfig -storage_server sts_hostname -stype PureDisk</pre>
4	Reuse or create storage path	<pre>/usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig -storage_server sts_hostname -stype PureDisk -configlist /root/local-lsu.txt</pre>

Reuse cloud LSU

Step	Task	Procedure
1	Reuse cloud LSU configuration.	<p>For each cloud LSU:</p> <pre>nbdevconfig -setconfig -storage_server sts_hostname -stype PureDisk -configlist /path/to/dr-lsu.txt</pre> <p>Config template example 1:</p> <pre>V7.5 "operation" "reuse-lsu-cloud" string V7.5 "lsuCloudUser" "XXX" string V7.5 "lsuCloudPassword" "XXX" string V7.5 "lsuCloudAlias" "<storageserver_lsuname>" string V7.5 "lsuCloudBucketName" "XXX" string V7.5 "lsuCloudBucketSubName" "XXX" string</pre> <p>Configuration template example 2 with encryption enabled:</p> <pre>V7.5 "operation" "reuse-lsu-cloud" string V7.5 "lsuCloudUser" "XXX" string V7.5 "lsuCloudPassword" "XXX" string V7.5 "lsuCloudAlias" "<storageserver_lsuname>" string V7.5 "lsuCloudBucketName" "XXX" string V7.5 "lsuCloudBucketSubName" "XXX" string V7.5 "lsuKmsServerName" "XXX" string</pre> <p>If alias does not exist, you can use the <code>csconfig</code> command to add them.</p> <pre>/usr/opensv/netbackup/bin/admincmd/csconfig cldinstance -as -in amazon.com -sts <storageserver> -lsu_name <lsuname></pre>
2	Recover <code>spad/spoold</code> metadata from cloud.	<p>For each cloud LSU:</p> <pre>/usr/opensv/pdde/pdcr/bin/cacontrol --catalog clouddr lsu1</pre>
3	Restart storage server.	<pre>/usr/opensv/netbackup/bin/bp.kill_all /usr/opensv/netbackup/bin/bp.start_all</pre> <p>Warning: The time taken can be high if the size of the containers is large. Use the following command to get the status:</p> <pre>/usr/opensv/pdde/pdcr/bin/cacontrol --catalog clouddrstatus lsu</pre>

Step	Task	Procedure
4	Start MSDP online check to recreate <code>refdb</code> .	<pre> /usr/openv/pdde/pdcr/bin/pddecfg -a enabledataintegritycheck /usr/openv/pdde/pdcr/bin/pddecfg -a startdatafullcheck /usr/openv/pdde/pdcr/bin/crcontrol --processqueue </pre>

About Image Sharing using MSDP cloud

Image Sharing lets you use an easy and visualized way to manage and provision images in cloud object storage and even the ability to convert backed up VMs as AWS instances in certain scenarios. Image sharing also supports MSDP cloud.

Before you begin, see the following topics:

- *Things to consider before you use Image sharing:* See [“About image sharing in cloud using Cloud Catalyst”](#) on page 289.
- *Using Image sharing:* See [“About image sharing in cloud using Cloud Catalyst”](#) on page 289.
- *Additional information about Image sharing:* See [“About image sharing in cloud using Cloud Catalyst”](#) on page 289.

Configure Image sharing using MSDP cloud

After installing NetBackup, you can run the `ims_system_config.py` script to configure Image sharing.

The path to access the command is: `/usr/openv/pdde/pdag/scripts/`

Amazon Web Service cloud provider

Use the following command:

```

ims_system_config.py -t PureDisk -k <key_id> -s <secret_key> -b
<bucket_name> -bs <bucket_sub_name> [-r <bucket_region>] [-p
<mount_point>]

```

Microsoft Azure cloud provider

Use the following command:

```

ims_system_config.py -cp 2 -t PureDisk -k <key_id> -s <secret_key>
-b <container_name> -bs <bucket_sub_name> [-p <mount_point>]

```

Other S3 compatible cloud provider (For example, Hitachi HCP)

Use the following command:

```
# ims_system_config.py -cp 3 -t PureDisk -k <key_id> -s <secret_key>  
-b <bucket_name> -bs <bucket_sub_name> -pt <cloud_provider_type> -sh  
<s3_hostname> -sp <s3_http_port> -sps <s3_https_port> -ssl <ssl_usage>
```

Note: Do not use `nbdevconfig` to modify cloud LSU or add new cloud LSU in image sharing server, that might cause an issue in the image sharing server.

If Amazon Glacier or deep Archive is enabled in MSDP cloud, you cannot use Image sharing.

As there is a requirement of free disk spaces by cloud LSU, ensure that you have enough disk space in the default mount point `/storage`, or you can use `-p` parameter of `ims_system_config.py` to specify a different mount point to meet the requirement of free disk spaces.

Monitoring deduplication activity

This chapter includes the following topics:

- [Monitoring the MSDP deduplication and compression rates](#)
- [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 and compression rates

The deduplication rate is the percentage of data that was stored by the deduplication engine. That data is not stored again. The compression rate is the percentage of space that is saved by compressing the backup data before it is stored.

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](#)


For the method to show the MSDP compression rate, See [“Viewing MSDP job details”](#) on page 327.

On UNIX and Linux, you can use the NetBackup `bpdjobs` command to display the deduplication rate. However, you must configure it to do so.

See [“To configure the `bpdjobs` command to display the MSDP deduplication rate”](#) on page 327.

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.

Not all columns appear by default. Click **View > Column Layout** to show or hide columns.

To configure the `bpdbjobs` command to display the MSDP deduplication rate

- ◆ Add a `DEDUPRATIO BPDBJOBS_COLDEFS` entry in the `bp.conf` file on the media server on which you run the command.

The `bpdbjobs` command then shows the deduplication rate in its output.

Disable the display of separate deduplication and compression rates

To disable the display of the compression rate separately:

- Open the `pd.conf` file that is available at the following locations:

Windows

```
<install_location>\lib\ost-plugins\pd.conf
```

UNIX

```
/usr/opensv/lib/ost-plugins/pd.conf
```

- Add the following parameter in the file:

```
DISPLAY_COMPRESSION_SPACE_SAVING = 0
```

You can remove this parameter or change the value to `1` to turn on the display of compression rate as a separate value.

Many factors affect deduplication performance.

See [“About MSDP performance”](#) on page 38.

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 328.

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.example.com deduplicated the data (the **dedup** field shows the deduplication rate and the **compression** field shows the storage space that is saved by compression):

```
LOG 1551428319 4 Info MSDP_Server.example.com 27726
StorageServer=PureDisk:MSDP_Server.example.com; Report=PDDO Stats
(multi-threaded stream used) for (MSDP_Server.example.com):
    scanned: 105098346 KB, CR sent: 2095410 KB, CR sent over FC: 0 KB,
    dedup: 98.0%, cache hits: 337282 (41.0%), where dedup space saving:89.7%,
    compression space saving:8.3%
```

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.

Additionally, for the client-side deduplication, the first **Info** line now displays the **dedupe** and **compression** values separately

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.example.com(pid=2220)
    Using OpenStorage client direct to backup from client
    Client_B.example.com to MSDP_Server.example.com
```

```
1/8/2013 11:58:09 PM - Info MSDP_Server.example.com(pid=2220)
StorageServer=PureDisk:MSDP_Server.example.com; Report=PDDO
Stats for (MSDP_Server.example.com: scanned: 110028 KB,
CR sent: 16654 KB, CR sent over FC: 0 KB, dedup: 84.9%,
cache disabled, where dedup space saving:3.8%,
compression space saving:81.1%
1/8/2013 11:58:09 PM - Info MSDP_Server.example.com(pid=2220)
Using the media server to write NBU data for backup
Client_B_1254987197.example.com to MSDP_Server.example.com
1/8/2013 11:58:19 PM - Info MSDP_Server.example.com(pid=2220)
StorageServer=PureDisk:MSDP_Server.example.com; Report=PDDO
Stats for (MSDP_Server.example.com: scanned: 17161 KB,
CR sent: 17170 KB, dedup: 0.0%, cache hits: 0 (0.0%)
```

Field descriptions

Table 8-1 describes the deduplication activity fields.

Table 8-1 MSDP activity field descriptions

Field	Description
Dedup space saving	The percentage of space that is saved by data deduplication (data is not written again).
Compression space saving	The percentage of space that is saved because the deduplication engine compressed some data before writing it to storage.
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 <code>pd.conf</code> file parameters” on page 210.</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 deduplication 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 the data to a NetBackup 5000 series appliance for storage. (In NetBackup, a PureDisk Storage Pool represents the storage on a NetBackup 5000 series appliance.)

Table 8-1 MSDP activity field descriptions (*continued*)

Field	Description
CR sent over FC	The amount of data that is sent from the deduplication plug-in over Fibre Channel to the component that stores the data. In NetBackup, the NetBackup Deduplication Engine stores the data.
dedup	The percentage of data that was stored already. That data is not stored again.
multi-threaded stream used	Indicates that the Deduplication Multi-Threaded Agent processed the backup. See “About the MSDP Deduplication Multi-Threaded Agent” on page 59.
PDDO stats	Indicates that the job details are for storage on the following destinations: <ul style="list-style-type: none">■ Media Server Deduplication Pool■ A NetBackup 5000 series appliance, exposed to NetBackup through a PureDisk Deduplication Pool.
rebased	The percentage of segments that were rebased (that is, defragmented) during the backup. Those segments had poor data locality. NetBackup reports backup job completion only after backup rebasing is completed. See “About MSDP storage rebasing” on page 368.
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 8-2](#) describes the options for monitoring capacity and usage.

Table 8-2 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.</p> <p>See “Monitoring the MSDP deduplication and compression rates” on page 325.</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 332.</p> <p>See “Viewing storage usage within MSDP container files” on page 332.</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 334.</p>

Table 8-2 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: 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 334.</p> <p>See “MSDP event codes and messages” on page 431.</p>
The nbdevquery 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 359.</p>
NetBackup OpsCenter	<p>The NetBackup OpsCenter also provides information about storage capacity and usage.</p> <p>See <i>NetBackup OpsCenter Administrator's Guide</i>: http://www.veritas.com/docs/DOC5332</p>

About MSDP container files

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

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

Viewing storage usage within MSDP container files

The NetBackup `crcontrol` command reports on storage usage within containers.

To view storage usage within MSDP container files

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

The following is an example of the command usage:

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

The following is an example of the output:

```
***** Data Store statistics *****
Data storage      Raw    Size   Used   Avail  Use%
                  1.0T  988.9G 666.0G 322.9G 68%

Number of containers      : 2981
Average container size    : 219740494 bytes (209.56MB)
Space allocated for containers : 655046415189 bytes (610.06GB)
Reserved space           : 45360705536 bytes (42.25GB)
Reserved space percentage : 4.1%
```

For systems that host a **Media Server Deduplication Pool**, you can use the following `crcontrol` command to show information about each partition:

```
/usr/opensv/pdde/pdcr/bin/crcontrol --dsstat 3
```

From the command output, you can determine the following:

Raw	The raw size of the storage.
Size	<p>The size of the storage that NetBackup can use: the Raw size of the storage minus the file system Reserved space.</p> <p>If the file system has a concept of root reserved space (such as EXT3 or VxFS), that space cannot be used for storage. The <code>crcontrol</code> command does not include reserved space in the available space. Unlike the <code>crcontrol</code> command, some operating system tools report root reserved space as usable space.</p>
Used	The amount of deduplicated data that is stored on the file system. NetBackup obtains the file system used space from the operating system.
Avail	The Size minus the Used space.
Use%	The Used space divided by the Size.

Viewing MSDP disk reports

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

[Table 8-3](#) describes the disk reports available.

Table 8-3 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 363.</p> <p>Either <code>PureDisk</code> or <code>Veritas 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 and Veritas Backup Exec are Veritas 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. This report displays only when a license that enables NetBackup disk features is installed.</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 390.

Table 8-4 Where to monitor the main MSDP processes

What	Where to monitor it
NetBackup Deduplication Engine	In the NetBackup Administration Console , the NetBackup Deduplication Engine appears as <code>spoold</code> on the Daemons tab of the Activity Monitor . The NetBackup <code>bpps</code> command also shows the <code>spoold</code> process.
NetBackup Deduplication Manager	In the NetBackup Administration Console , the NetBackup Deduplication Manager appears as <code>spad</code> on the Daemons tab of the Activity Monitor . The NetBackup <code>bpps</code> command also shows the <code>spad</code> process.

Reporting on Auto Image Replication jobs

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

Table 8-5 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 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](#)
- [How MSDP restores work](#)
- [Configuring MSDP restores directly to a client](#)
- [About restoring files at a remote site](#)
- [About restoring from a backup at a target master domain](#)
- [Specifying the restore server](#)

Managing MSDP servers

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

See [“Viewing MSDP storage servers”](#) on page 338.

See [“Determining the MSDP storage server state”](#) on page 338.

See [“Viewing MSDP storage server attributes”](#) on page 339.

See [“Setting MSDP storage server attributes”](#) on page 340.

See [“Changing MSDP storage server properties”](#) on page 341.

See [“Clearing MSDP storage server attributes”](#) on page 342.

See [“About changing the MSDP storage server name or storage path”](#) on page 343.

See [“Changing the MSDP storage server name or storage path”](#) on page 343.

See [“Removing an MSDP load balancing server”](#) on page 345.

See [“Deleting an MSDP storage server”](#) on page 346.

See [“Deleting the MSDP storage server configuration”](#) on page 347.

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.example.com
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 352.

To set a MSDP storage server attribute

- 1 The following is the command syntax to set a storage server attribute. Run the command on the master server or on the storage server.

```
nbdevconfig -changests -storage_server storage_server -stype  
PureDisk -setattribute attribute
```

The following describes the options that require the arguments that are specific to your domain:

<code>-storage_server</code>	The name of the storage server.
<code>storage_server</code>	
<code>-setattribute</code>	The <i>attribute</i> is the name of the argument that represents the new functionality.
<code>attribute</code>	

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 339.

See [“About MSDP optimized synthetic backups”](#) on page 36.

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 Servers**.
- 2** Select the deduplication storage server.

Warning: If you have load balancing servers, do not select any or all of them in the **Media Servers** pane of the **NetBackup Administration Console**. If you do, the change operation fails.

- 3** On the **Edit** menu, select **Change**.

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



- 5 For the property to change, select the value in the **Value** column.
- 6 Change the value.
- 7 Click **OK**.

Clearing MSDP storage server attributes

Use the `nbdevconfig` command to remove storage server attributes.

To clear MSDP storage server attributes

- ◆ Run the following command on the NetBackup master server or on a storage server:

```
nbdevconfig -changests -storage_server storage_server -stype
PureDisk -clearattribute attribute
```

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

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

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

- UNIX: `/usr/openv/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 343.
- 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 343.
- 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 381.

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 343.

Table 9-1 Changing the storage server name or storage path

Step	Task	Procedure
Step 1	Ensure that no deduplication activity occurs	Deactivate all backup policies that use deduplication storage. See the <i>NetBackup Administrator's Guide, Volume I</i> : http://www.veritas.com/docs/DOC5332
Step 2	Expire the backup images	Expire all backup images that reside on the deduplication disk storage. Warning: Do not delete the images. They are imported back into NetBackup later in this process. If you use the <code>bpexpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter. See the <i>NetBackup Administrator's Guide, Volume I</i> : http://www.veritas.com/docs/DOC5332
Step 3	Delete the storage units that use the disk pool	See the <i>NetBackup Administrator's Guide, Volume I</i> : http://www.veritas.com/docs/DOC5332
Step 4	Delete the disk pool	See "Deleting a Media Server Deduplication Pool" on page 361.
Step 5	Delete the deduplication storage server	See "Deleting an MSDP storage server" on page 346.
Step 6	Delete the configuration	Delete the deduplication configuration. See "Deleting the MSDP storage server configuration" on page 347.
Step 7	Delete the deduplication host configuration file	Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers. See "Deleting an MSDP host configuration file" on page 232.
Step 8	Delete the identity file and the file system table file	Delete the following files from the MSDP storage server, depending on operating system: UNIX: <code>/storage_path/data/.identity</code> <code>/storage_path/etc/puredisk/fstab.cfg</code> Windows: <code>storage_path\data\.identity</code> <code>storage_path\etc\puredisk\fstab.cfg</code>

Table 9-1 Changing the storage server name or storage path (*continued*)

Step	Task	Procedure
Step 9	Change the storage server name or the storage location	See the computer or the storage vendor's documentation. See "Use fully qualified domain names" on page 41. See "MSDP storage path properties" on page 94.
Step 10	Reconfigure the storage server	When you configure deduplication, select the host by the new name and enter the new storage path (if you changed the path). You can also use a new network interface. See "Configuring MSDP server-side deduplication" on page 55.
Step 11	Import the backup images	See the <i>NetBackup Administrator's Guide, Volume I</i> : http://www.veritas.com/docs/DOC5332

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 27.

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

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

For procedures, see *NetBackup Administrator's Guide, Volume II*:


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

To remove a media server from a MSDP node

- 1 For every storage unit that specifies the media server in **Use one of the following media servers**, clear the check box that specifies the media server.

This step is not required if the storage unit is configured to use any available media server.
- 2 In the NetBackup Administration Console, expand **Media and Device Management > Credentials > Storage Server**.

- 3 Select the deduplication storage server, and 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 343.

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 343.

See [“Deactivating MSDP”](#) on page 388.

To delete the MSDP storage server configuration

- 1 Use the NetBackup Administration Console to stop the NetBackup Deduplication Engine (`spoold`) and the NetBackup Deduplication Manager (`spad`).
- 2 On the storage server, run one of the following scripts, depending on your operating system:

UNIX:

```
/usr/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.

Managing NetBackup Deduplication Engine credentials

You can manage existing credentials in NetBackup.

See [“Determining which media servers have deduplication credentials”](#) on page 348.

See [“Adding NetBackup Deduplication Engine credentials”](#) on page 348.

See [“Changing NetBackup Deduplication Engine credentials”](#) on page 349.

See [“Deleting credentials from a load balancing server”](#) on page 349.

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, and 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 206.

To add NetBackup Deduplication Engine credentials by using the `tpconfig` command

- ◆ On the host to which you want to add credentials, run the following command:

On Windows:

```
install_path\Veritas\NetBackup\Volmgr\bin\tpconfig -add
-storage_server sshostname -stype PureDisk -sts_user_id UserID
-password Password
```

On UNIX/Linux:

```
/usr/opensv/volmgr/bin/tpconfig -add -storage_server sshostname
-stype PureDisk -sts_user_id UserID -password Password
```

For *sshostname*, use the name of the storage server.

Changing NetBackup Deduplication Engine credentials

You cannot change the NetBackup Deduplication Engine credentials after you enter them. If you must change the credentials, contact your Veritas support representative.

See [“About the NetBackup Deduplication Engine credentials”](#) on page 34.

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 345.

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/opensv/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 350.

See [“Changing a Media Server Deduplication Pool properties”](#) on page 353.

See [“Determining the Media Server Deduplication Pool state”](#) on page 350.

See [“Changing Media Server Deduplication Pool state”](#) on page 351.

See [“Determining the MSDP disk volume state”](#) on page 359.

See [“Changing the MSDP disk volume state”](#) on page 359.

See [“Viewing Media Server Deduplication Pool attributes”](#) on page 351.

See [“Setting a Media Server Deduplication Pool attribute”](#) on page 352.

See [“Clearing a Media Server Deduplication Pool attribute”](#) on page 358.

See [“Resizing the MSDP storage partition”](#) on page 372.

See [“Inventorying a NetBackup disk pool”](#) on page 360.

See [“Deleting a Media Server Deduplication Pool”](#) on page 361.

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 Media Server Deduplication Pool state

You can change the state of a NetBackup disk pool. Pool states are 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 Media Server Deduplication Pool state

- 1 In the **NetBackup Administration Console**, in the left pane, select **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**.

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.example.com (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 340.

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 351.


Changing a Media Server Deduplication Pool properties

You can change the properties of a deduplication disk pool.

To change disk pool properties

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices > Disk Pools**.
- 2 Select the disk pool you want to change in the details pane.

- 3 On the **Edit** menu, select **Change**.



Change Disk Pool

Name: dp-msdp-ss-host

Storage servers: (PureDisk) msdp-ss-host.example.com

Disk volumes:

Volume Name	Available Space	Raw Size	Replication
PureDiskVolume	22.17 GB	22.25 GB	None

Refresh

Total raw size: 22.25 GB
 Total available space: 22.17 GB
 Targeted replication: No

Comments:

Disk Volume Settings

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

Maximum I/O Streams

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

☐ Limit I/O streams: 2 per volume

OK Cancel Help

- 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 355.
- 5 Change the other properties as necessary.
 See [“Media Server Deduplication Pool properties”](#) on page 116.
- 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.

It is recommended that you take the following actions when the volume topology changes:

- 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 9-2 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 9-2 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 9-2 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 165.</p> <p>Work with your storage administrator to understand the changes and why they were made. The replication relationships may or may not have to be re-established. If the relationship was removed in error, re-establishing the relationships seem justified. If you are retiring or replacing the target replication device, you probably do not want to re-establish the relationships.</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 9-2 Refresh outcomes (*continued*)

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

Clearing 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> <i>pool_name</i>	The name of the disk pool.
<code>-setattribute</code> <i>attribute</i>	The <i>attribute</i> is the name of the argument that represents the new functionality.

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

- UNIX: `/usr/openv/netbackup/bin/admincmd`
- Windows: `install_path\NetBackup\bin\admincmd`

Determining the MSDP disk volume state

Use the NetBackup `nbdevquery` command to determine the state of the volume in a deduplication disk pool. NetBackup exposes all of the storage for MSDP as a single volume, **PureDiskVolume**. The command shows the properties and attributes of the **PureDiskVolume**.

To determine MSDP disk volume state

- ◆ Display the volume state by using the following command:

UNIX: `/usr/openv/netbackup/bin/admincmd/nbdevquery -listdv -stype PureDisk -U -dp disk_pool_name`

Windows: `install_path\NetBackup\bin\admincmd\nbdevquery -listdv -stype PureDisk -U -dp disk_pool_name`

The *state* is either UP or DOWN.

The following is example output

```
Disk Pool Name      : MSDP_Disk_Pool
Disk Type           : PureDisk
Disk Volume Name    : PureDiskVolume
Disk Media ID       : @aaaab
Total Capacity (GB) : 49.98
Free Space (GB)     : 43.66
Use%                : 12
Status              : UP
Flag                : ReadOnWrite
Flag                : AdminUp
Flag                : InternalUp
Num Read Mounts     : 0
Num Write Mounts    : 1
Cur Read Streams   : 0
Cur Write Streams  : 0
```

Changing the MSDP disk volume state

The disk volume state is **UP** or **DOWN**. NetBackup exposes all of the storage for MSDP as a single volume, **PureDiskVolume**.

To change the state to **DOWN**, the disk pool in which the volume resides must not be busy. If backup jobs are assigned to the disk pool, the state change fails. Cancel the backup jobs or wait until the jobs complete.

To change the MSDP disk volume state

- ◆ Change the disk volume state; the following is the command syntax:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbdevconfig -changestate
-stype PureDisk -dp disk_pool_name -dv PureDiskVolume -state state`

Windows: `install_path\NetBackup\bin\admincmd\nbdevconfig
-changestate -stype PureDisk -dp disk_pool_name -dv PureDiskVolume
-state state`

For the `-state`, specify either **UP** or **DOWN**.

Inventorying a NetBackup disk pool

An inventory of a NetBackup disk pool reads the capacity of the disk volumes in the pool. An inventory operation lets you update NetBackup with the new capacity values if you do the following:

- Increase or decrease the size of the disk volumes in a disk pool.
- Add volumes to or remove volumes from a disk pool.

How you increase or decrease the underlying storage capacity depends on your storage implementation. You must complete that process before you inventory the disk pool.

To inventory a NetBackup disk pool

- 1 in the **NetBackup Administration Console**, select **Media and Device Management > Devices > Disk Pools**.
- 2 On the **Actions** menu, select **Inventory Disk Pools**.
- 3 In the **Inventory Disk Pool** dialog box, select the disk pool to inventory and then click **Start Inventory**.
- 4 To update the NetBackup catalog with the capacity values returned by the inventory, click **Update Configuration**.
- 5 To inventory another disk pool, go to step 3.
- 6 To exit, click **Close**.

Deleting a Media Server Deduplication Pool

You can delete a disk pool if it does not contain valid NetBackup backup images or image fragments. If it does, you must first expire and delete those images or fragments. If expired image fragments remain on disk, you must remove those also.

See [“Cannot delete an MSDP disk pool”](#) on page 422.

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, Veritas recommends the following approach.

See [“About the MSDP data removal process”](#) on page 371.

To delete backup images manually

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

Without this option, `bpexpdate` calls `nbdelete` to delete images. Each call to `nbdelete` creates a job in the Activity Monitor, allocates resources, and launches processes on the media server.

- 2 After you expire the last image, delete all of the images by using the `nbdelete` command with the `-allvolumes` option.

Only one job is created in the Activity Monitor, fewer resources are allocated, and fewer processes are started on the media servers. The entire process of expiring images and deleting images takes less time.

About MSDP queue processing

Operations that require database updates accumulate in a transaction queue. Twice a day, the NetBackup Deduplication Manager directs the Deduplication Engine to process the queue as one batch. By default, queue processing occurs every 12 hours, 20 minutes past the hour.

Primarily, the transaction queue contains clean-up and integrity checking transactions. These transactions update the reference database.

Queue processing writes status information to the deduplication engine `stored.log` file.

See [“NetBackup MSDP log files”](#) on page 407.

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 Veritas 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 362.

See [“About MSDP server requirements”](#) on page 28.

Processing the MSDP transaction queue manually

NetBackup maintains a queue for MSDP database transactions.

See [“About MSDP queue processing”](#) on page 362.

Usually, you should not need to run the deduplication database transaction queue processes manually. However, when you recover the MSDP catalog from a backup, you must process the MSDP transaction queue. Processing the transaction queue is part of a larger process.

To process the MSDP transaction queue manually

- 1 On the MSDP storage server, run the following command:

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

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

- 2 To determine if the queue processing is still active, run the following command:

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

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

`--processqueueinfo`

If the output shows `Busy : yes`, the queue is still active.

- 3 To examine the results, run the following command (number 1 not lowercase letter l):

UNIX: `/usr/opensv/pdde/pdcr/bin/crcontrol --dsstat 1`

Windows: `install_path\Veritas\pdde\Crcontrol.exe --dsstat 1`

The command may run for a long time; if you omit the `1`, results return more quickly but they are not as accurate.

About MSDP data integrity checking

Deduplication metadata and data may become inconsistent or corrupted because of disk failures, I/O errors, database corruption, and operational errors. NetBackup checks the integrity of the deduplicated data on a regular basis. NetBackup performs some of the integrity checking when the storage server is idle. Other integrity checking is designed to use few storage server resources so as not to interfere with operations.

The data integrity checking process includes the following checks and actions:

- Automatically constrains data loss or corruption to ensure that new backups are intact.
- Automatically runs a cyclic redundancy check (CRC) for the data containers.
- Automatically collects and cleans up storage garbage.
- Automatically recovers the container-based reference database (or parts of the database) if it is corrupt or missing.
- Automatically finds storage leaks and fixes them.

NetBackup resolves many integrity issues without user intervention, and some issues are fixed when the next backup runs. However, a severe issue may require

intervention by Veritas Support. In such cases, NetBackup writes a message to the NetBackup Disk Logs report.

See [“Viewing MSDP disk reports”](#) on page 334.

The data integrity message code is 1057.

See [“MSDP event codes and messages”](#) on page 431.

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

See [“NetBackup MSDP log files”](#) on page 407.

You can configure some of the data integrity checking behaviors.

See [“Configuring MSDP data integrity checking behavior”](#) on page 364.

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 364.
- Edit configuration file parameters.
See [“To configure data integrity checking behavior by editing the configuration files”](#) on page 366.

Warning: Veritas 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 363.

See [“MSDP data integrity checking configuration parameters”](#) on page 366.

To configure data integrity checking behavior by using a command

- ◆ To configure behavior, specify a value for each of the data integrity checks, as follows:
 - Data consistency checking. Use the following commands to configure behavior:

Enable	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a enabledataintegritycheck</code></p> <p>Windows: <code>install_path\Veritas\pdde\pddecfg -a enabledataintegritycheck</code></p>
Disable	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a disabledataintegritycheck</code></p> <p>Windows: <code>install_path\Veritas\pdde\pddecfg -a disabledataintegritycheck</code></p>
Get the status	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/pddecfg -a getdataintegritycheck</code></p> <p>Windows: <code>install_path\Veritas\pdde\pddecfg -a getdataintegritycheck</code></p>

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

Enable	<p>CRC does not run if queue processing is active or during disk read or write operations.</p> <p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckon</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --crccheckon</code></p>
Disable	<p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckoff</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --crccheckoff</code></p>
Enable fast checking	<p>Fast check CRC mode begins the check from container 64 and does not sleep between checking containers.</p> <p>When the fast CRC ends, CRC behavior reverts to the behavior before fast checking was invoked.</p> <p>UNIX: <code>/usr/opensv/pdde/pdcr/bin/crcontrol --crccheckrestart</code></p> <p>Windows: <code>install_path\Veritas\pdde\Crcontrol.exe --crccheckrestart</code></p>

```
Get the status  UNIX: /usr/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 366.

- 3 Save and close the file.
- 4 Restart the NetBackup Deduplication Engine and the NetBackup Deduplication Manager.

You can do this from the **Daemons** tab in the **Activity Monitor**.

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 9-3](#).
See “[About the MSDP contentrouter.cfg file](#)” on page 227.
- The `spa.cfg` file.
The parameters are described in [Table 9-3](#).

Those files reside in the following directories:

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

Warning: Veritas 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 363.

Table 9-3 The contentrouter.cfg file parameters for data integrity checking

Setting	Default	Description
EnableCRCCheck	true	<p>Enable or disable cyclic redundancy checking (CRC) of the data container files.</p> <p>The possible values are <code>true</code> or <code>false</code>.</p> <p>CRC occurs only when no backup, restore, or queue processing jobs are running.</p>
CRCCheckSleepSeconds	5	<p>The time in seconds to sleep between checking containers.</p> <p>The longer the sleep interval, the more time it takes to check containers.</p>
CRCCheckBatchNum	40	<p>The number of containers to check each time.</p> <p>The greater the number of containers, the less time it takes to check all containers, but the more system resources it takes.</p>
ShutdownCRWhenError	false	<p>Stops the NetBackup Deduplication Manager when a data loss is discovered.</p> <p>This parameter is reserved for debugging purposes by Veritas Support Representatives.</p> <p>The possible values are <code>true</code> or <code>false</code>.</p>
GarbageCheckRemainDCCount	100	<p>The number of containers from failed jobs not to check for garbage. A failed backup or replication job still produces data containers. Because failed jobs are retried, retaining those containers means NetBackup does not have to send the fingerprint information again. As a result, retried jobs consume less time and fewer system resources than when first run.</p>

Table 9-4 spa.cfg file parameters for data integrity checking

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

Table 9-4 spa.cfg file parameters for data integrity checking (*continued*)

Setting	Default	Description
EnableDataCheckAlert	true	<p>Enable or disable alerts.</p> <p>If <code>true</code>, NetBackup writes a message to the Disk Logs report when it detects a lost data segment.</p> <p>See “NetBackup MSDP log files” on page 407.</p>

About managing MSDP storage read performance

NetBackup provides some control over the processes that are used for read operations. The read operation controls can improve performance for the jobs that read from the storage. Such jobs include restore jobs, duplication jobs, and replication jobs.

In most cases, you should change configuration file options only when directed to do so by a Veritas support representative.

Defragment the storage

NetBackup includes a process, called *rebasing*, which defragments the backup images in a deduplication pool. Read performance improves when the file segments from a client backup are close to each other on deduplication storage.

See [“About MSDP storage rebasing”](#) on page 368.

Decrypt the data on the client rather than the server

The `RESTORE_DECRYPT_LOCAL` parameter in the `pd.conf` file specifies on which host to decrypt and decompress the data during restore operations.

See [“About the MSDP pd.conf configuration file”](#) on page 209.

See [“MSDP pd.conf file parameters”](#) on page 210.

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.

After the rebasing, NetBackup reclaims the storage space that the rebased data segments used.

[Table 9-5](#) describes the rebasing operations.

Table 9-5 Types of rebasing

Type	Description
Normal backup rebasing	<p>The rebasing that occurs during a backup if the normal rebasing criteria are met, as follows:</p> <ul style="list-style-type: none">■ The container has been rebased within the last 3 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. <p>See “MSDP pd.conf file parameters” on page 210.</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> <p>NetBackup reports backup job completion after the rebasing is completed.</p>

Table 9-5 Types of rebasing (*continued*)

Type	Description
Periodic backup rebasing	<p>The rebasing that occurs during a backup if the periodic rebasing criteria are met, as follows:</p> <ul style="list-style-type: none">■ The container has not been rebased within the last 3 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 pd.conf file parameters” on page 210.</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> <p>NetBackup reports backup job completion after the rebasing is completed.</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>Some parameters in the <code>contentrouter.cfg</code> file control the server-side rebasing behavior.</p> <p>See “MSDP server-side rebasing parameters” on page 370.</p>

MSDP server-side rebasing parameters

[Table 9-6](#) describes the parameters that control server-side rebasing.

See [“About MSDP storage rebasing”](#) on page 368.

Usually, you do not need to change parameter values. However, in some cases, you may be directed to change settings by a Veritas support representative.

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

See [“About the MSDP contentrouter.cfg file”](#) on page 227.

Table 9-6 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 332.

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 363.

See “[Deleting backup images](#)” on page 361.

Resizing the MSDP storage partition

If the volume that contains the deduplication storage is resized dynamically, restart the NetBackup services on the storage server. You must restart the services so that NetBackup can use the resized partition correctly. If you do not restart the services, NetBackup reports the capacity as full prematurely.

To resize the MSDP storage

- 1 Stop all NetBackup jobs on the storage on which you want to change the disk partition sizes and wait for the jobs to end.

- 2 Deactivate the media server that hosts the storage server.

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

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

- 3 Stop the NetBackup services on the storage server.

Be sure to wait for all services to stop.

- 4 Use the operating system or disk manager tools to dynamically increase or decrease the deduplication storage area.

- 5 Restart the NetBackup services.

- 6 Activate the media server that hosts the storage server.

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

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

- 7 Restart the deduplication jobs.

See “[About adding additional MSDP storage](#)” on page 49.

How MSDP restores work

The following two methods exist to for MSDP restore operations:

Table 9-7 MSDP restore types

Type	Description
Normal restore	<p>The MSDP storage server first <i>rehydrates</i> (that is, reassembles) the data. NetBackup then chooses the least busy media server to move the data to the client. (NetBackup chooses the least busy media server from those that have credentials for the NetBackup Deduplication Engine.) The media server <code>bptm</code> process moves the data to the client.</p> <p>The following media servers have credentials for the NetBackup Deduplication Engine:</p> <ul style="list-style-type: none">■ The media server that hosts the storage server. Although the media server and the storage server share a host, the storage server sends the data through the media server <code>bptm</code> process on that host.■ A load balancing server in the same deduplication node. See “About MSDP load balancing servers” on page 28.■ A deduplication server in a different deduplication node that is the target of optimized duplication. See “About the media servers for MSDP optimized duplication within the same domain” on page 144. <p>You can specify the server to use for restores.</p> <p>See “Specifying the restore server” on page 375.</p>
Restore directly to the client	<p>The storage server can bypass the media server and move the data directly to the client. NetBackup does not choose a media server for the restore, and the restore does not use the media server <code>bptm</code> process.</p> <p>You must configure NetBackup to bypass a media server and receive the restore data directly from the storage server.</p> <p>See “Configuring MSDP restores directly to a client” on page 374.</p> <p>By default, NetBackup rehydrates the data on the storage server except for client-side deduplication clients. Those clients rehydrate the data. You can configure NetBackup so that the data is rehydrated on the storage server rather than the client. See the <code>RESTORE_DECRYPT_LOCAL</code> parameter in the MSDP <code>pd.conf</code> file.</p> <p>See “MSDP pd.conf file parameters” on page 210.</p> <p>See “Editing the MSDP pd.conf file” on page 210.</p>

Configuring MSDP restores directly to a client

The NetBackup MSDP storage server can move restore data directly to an MSDP client, bypassing the media server components.

See “[How MSDP restores work](#)” on page 373.

To enable restores directly to a client

- 1 Set the `OLD_VNETD_CALLBACK` option to `YES` on the client. The `OLD_VNETD_CALLBACK` option is stored in the `bp.conf` file on UNIX systems and the registry on Windows systems.

See “[Setting NetBackup configuration options by using the command line](#)” on page 155.

- 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`

About restoring files at a remote site

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

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

See “Managing client restores” in the *NetBackup Administrator's Guide, Volume I*:
<http://www.veritas.com/docs/DOC5332>

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

About restoring from a backup at a target master domain

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

Table 9-8 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 373.

You can specify the server to use for restores. The following are the methods that specify the restore server:

- Always use the backup server. Two methods exist, as follows:
 - Use NetBackup **Host Properties** to specify a **Media host override** server. All restore jobs for any storage unit on the original backup server use the media server you specify. Specify the same server for the **Restore server** as for the **Original backup server**.

See “Forcing restores to use a specific server” in the *NetBackup Administrator’s Guide, Volume I*:

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

This procedure sets the `FORCE_RESTORE_MEDIA_SERVER` option. Configuration options are stored in the `bp.conf` file on UNIX systems and the registry on Windows systems.

- Create the touch file `USE_BACKUP_MEDIA_SERVER_FOR_RESTORE` on the NetBackup master server in the following directory:

UNIX: `usr/opensv/netbackup/db/config`

Windows: `install_path\Veritas\Netbackup\db\config`

This global setting always forces restores to the server that did the backup. It applies to all NetBackup restore jobs, not just deduplication restore jobs.

If this touch file exists, NetBackup ignores the `FORCE_RESTORE_MEDIA_SERVER` and `FAILOVER_RESTORE_MEDIA_SERVER` settings.

- Always use a different server.
Use NetBackup **Host Properties** to specify a **Media host override** server. See the previous explanation about **Media host override**, except: Specify the different server for the **Restore server**.
- A single restore instance. Use the `bprestore` command with the `-disk_media_server` option.
Restore jobs for each instance of the command use the media server you specify. See the *NetBackup Commands Reference Guide*:

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

Recovering MSDP

This chapter includes the following topics:

- [About recovering the MSDP catalog](#)
- [Restoring the MSDP catalog from a shadow copy](#)
- [Recovering from an MSDP storage server disk failure](#)
- [Recovering from an MSDP storage server failure](#)
- [Recovering the MSDP storage server after NetBackup catalog recovery](#)

About recovering the MSDP catalog

The following are the recovery options for the NetBackup MSDP catalog:

Table 10-1 MSDP catalog backup recovery options

Recovery option	Description
Restore from a shadow copy	<p>If NetBackup detects corruption in the MSDP catalog, the Deduplication Manager restores the catalog automatically from the most recent shadow copy. The automatic restore process also plays a transaction log so that the recovered MSDP catalog is current.</p> <p>Although the shadow copy restore process is automatic, a restore procedure is available if you need to recover from a shadow copy manually.</p> <p>See “About the MSDP shadow catalog” on page 233.</p> <p>See “Restoring the MSDP catalog from a shadow copy” on page 378.</p>

Table 10-1 MSDP catalog backup recovery options (continued)

Recovery option	Description
Recover from a backup	<p>If you configured an MSDP catalog backup policy and a valid backup exists, you can recover the catalog from a backup. As a general rule, you should only attempt to recover the MSDP catalog from a backup if you have no alternatives. As an example: A hardware problem or a software problem results in the complete loss of the MSDP catalog and the shadow copies.</p> <p>The greatest chance for a successful outcome when you recover the MSDP catalog from a backup is when the recovery is guided. An unsuccessful outcome may cause data loss. For the customers who need to recover the MSDP catalog, Veritas wants to guide them through the process. Therefore, to recover the MSDP catalog from a backup, contact your Veritas support representative. You can refer the support representative to Knowledge Base Article 000047346, which contains the recovery instructions.</p>

Caution: You must determine if your situation is severe enough to recover the catalog. Veritas recommends that you contact your Veritas Support representative before you restore or recover the MSDP catalog. The Support representative can help you determine if you need to recover the catalog or if other solutions are available.

See [“About protecting the MSDP catalog”](#) on page 233.

Restoring the MSDP catalog from a shadow copy

NetBackup automatically restores the necessary parts of the MSDP catalog if corruption is detected. However, you can restore the MSDP catalog from a shadow copy manually, although in normal circumstances it is not necessary. Veritas recommends that you contact your Veritas Support representative before you restore all or part of the MSDP catalog from a shadow copy.

The procedure that you use depends on the restore scenario, as follows:

- Restore the entire MSDP catalog from a shadow copy

In this scenario, you want to restore the entire catalog from one of the shadow copies.

See [“To restore the entire MSDP catalog from a shadow copy”](#) on page 379.

Restore a specific MSDP database file

The MSDP catalog is composed of multiple small database files. Those files are organized in the file system by the client name and policy name, as follows:

UNIX:

/database_path/databases/catalogshadow/2/ClientName/PolicyName

Windows:

database_path\databases\catalogshadow\2\ClientName\PolicyName

You can restore the database files for a client and a policy combination. The restore of a specific client's and policy's database files is always from the most recent shadow copy.

See [“To restore a specific MSDP database file from a shadow copy”](#) on page 379.

See [“About recovering the MSDP catalog”](#) on page 377.

To restore the entire MSDP catalog from a shadow copy

- 1 If any MSDP jobs are active, either cancel them or wait until they complete.
- 2 Disable all policies and storage lifecycle policies that back up to the **Media Server Deduplication Pool**.
- 3 On the MSDP storage server, run the following command, depending on host type:

- UNIX: `/usr/opensv/pdde/pdcr/bin/cacontrol --catalog recover all`
- Windows: `install_path\Veritas\pdde\cacontrol --catalog recover all`

- 4 Enable all policies and storage lifecycle policies that back up to the **Media Server Deduplication Pool**.

- 5 Restart the jobs that were canceled before the recovery.

To restore a specific MSDP database file from a shadow copy

- 1 If any MSDP jobs are active for the client and the backup policy combination, either cancel them or wait until they complete.
- 2 Disable the policies and storage lifecycle policies for the client and the backup policy combination that back up to the **Media Server Deduplication Pool**.

- 3 Change to the shadow directory for the client and policy from which you want to recover that database file. That directory contains the database files from which to recover. The following are the pathname formats:

UNIX:

/database_path/databases/catalogshadow/2/ClientName/PolicyName

Windows:

database_path\databases\catalogshadow\2\ClientName\PolicyName

- 4 Run the following command, depending on host type:
 - UNIX: `/usr/opensv/pdde/pdcr/bin/cacontrol --catalog recover 2 "/ClientName/PolicyName"`
 - Windows: `install_path\Veritas\pdde\cacontrol --catalog recover 2 "\ClientName\PolicyName"`
- 5 Enable all policies and storage lifecycle policies that back up to the **Media Server Deduplication Pool**.
- 6 If you canceled jobs before you recovered the database files, restart them.

Recovering from an MSDP storage server disk failure

If recovery mechanisms do not protect the disk on which the NetBackup software resides, the deduplication storage server configuration is lost if the disk fails. This topic describes how to recover from a system disk or program disk failure where the disk was not backed up.

Note: This procedure describes recovery of the disk on which the NetBackup media server software resides not the disk on which the deduplicated data resides. The disk may or may not be the system boot disk.

After recovery, your NetBackup deduplication environment should function normally. Any valid backup images on the deduplication storage should be available for restores.

Veritas 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 10-2 Process to recover from media server disk failure

Step	Task	Procedure
Step 1	Replace the disk.	If the disk is a system boot disk, also install the operating system. See the hardware vendor and operating system documentation.
Step 2	Mount the storage.	Ensure that the storage and database are mounted at the same locations. See the storage vendor's documentation.
Step 3	Install and license the NetBackup media server software.	See <i>NetBackup Installation Guide for UNIX and Windows</i> : http://www.veritas.com/docs/DOC5332 See "About the MSDP license" on page 51.
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 232.
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 349.
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 348.
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 228.
Step 8	Edit the configuration file	See "Editing an MSDP storage server configuration file" on page 229.
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 230.
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 206.

Recovering from an MSDP storage server failure

To recover from a permanent failure of the storage server host computer, use the process that is described in this topic.

NetBackup recommends that you consider the following items before you recover:

- The new computer must use the same byte order as the old computer.

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

- Veritas recommends that the new computer use the same operating system as the old computer.
- Veritas recommends that the new computer use the same version of NetBackup as the old computer.

If you use a newer version of NetBackup on the new computer, ensure that you perform any data conversions that may be required for the newer release.

If you want to use an older version of NetBackup on the replacement host, contact your Veritas support representative.

Table 10-3 Recover from an MSDP storage server failure

Step	Task	Procedure
Step 1	Expire the backup images	<p>Expire all backup images that reside on the deduplication disk storage.</p> <p>Warning: Do not delete the images. They are imported back into NetBackup later in this process.</p> <p>If you use the <code>bpxupdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>: http://www.veritas.com/docs/DOC5332</p>
Step 2	Delete the storage units that use the disk pool	<p>See the <i>NetBackup Administrator's Guide, Volume I</i>: http://www.veritas.com/docs/DOC5332</p>
Step 3	Delete the disk pool	See "Deleting a Media Server Deduplication Pool" on page 361.
Step 4	Delete the deduplication storage server	See "Deleting an MSDP storage server" on page 346.

Table 10-3 Recover from an MSDP storage server failure (*continued*)

Step	Task	Procedure
Step 5	Delete the deduplication host configuration file	<p>Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers.</p> <p>See “Deleting an MSDP host configuration file” on page 232.</p>
Step 6	Delete the credentials on deduplication servers	<p>If you have load balancing servers, delete the NetBackup Deduplication Engine credentials on those media servers.</p> <p>See “Deleting credentials from a load balancing server” on page 349.</p>
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 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. <p>If the original host used a specific network interface, you do not have to use the same interface name.</p> ■ If you had configured the previous MSDP storage server to use MSDP Encryption using KMS service, you must use the same configuration for the new MSDP storage server. <p>See “About MSDP storage servers” on page 27.</p> <p>See “About MSDP server requirements” on page 28.</p>
Step 8	Connect the storage to the host	<p>Use the storage path that you configured for this replacement host.</p> <p>See the computer or the storage vendor's documentation.</p>
Step 9	Install the NetBackup media server software on the new host	<p>See the <i>NetBackup Installation Guide for UNIX and Windows</i>: http://www.veritas.com/docs/DOC5332</p>

Table 10-3 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 server-side deduplication” on page 55.
Step 11	Import the backup images	See the <i>NetBackup Administrator’s Guide, Volume I</i> : http://www.veritas.com/docs/DOC5332

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 230.

Veritas recommends that you save your storage server configuration.

See [“Save the MSDP storage server configuration”](#) on page 46.

Information about recovering the master server is available.

See *NetBackup Troubleshooting Guide*:

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

Replacing MSDP hosts

This chapter includes the following topics:

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

Replacing the MSDP storage server host computer

If you replace the deduplication storage server host computer, use these instructions to install NetBackup and reconfigure the deduplication storage server. The new host cannot host a deduplication storage server already.

Reasons to replace the computer include a lease swap or perhaps the current deduplication storage server computer does not meet your performance requirements.

NetBackup recommends that you consider the following items before you recover:

- The new computer must use the same byte order as the old computer.

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

- Veritas recommends that the new computer use the same operating system as the old computer.
- Veritas recommends that the new computer use the same version of NetBackup as the old computer.

If you use a newer version of NetBackup on the new computer, ensure that you perform any data conversions that may be required for the newer release.

If you want to use an older version of NetBackup on the replacement host, contact your Veritas support representative.

Table 11-1 Replacing an MSDP storage server host computer

Step	Task	Procedure
Step 1	Expire the backup images	<p>Expire all backup images that reside on the deduplication disk storage.</p> <p>Warning: Do not delete the images. They are imported back into NetBackup later in this process.</p> <p>If you use the <code>bpexpdate</code> command to expire the backup images, use the <code>-nodelete</code> parameter.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>: http://www.veritas.com/docs/DOC5332</p>
Step 2	Delete the storage units that use the disk pool	<p>See the <i>NetBackup Administrator's Guide, Volume I</i>: http://www.veritas.com/docs/DOC5332</p>
Step 3	Delete the disk pool	See "Deleting a Media Server Deduplication Pool" on page 361.
Step 4	Delete the deduplication storage server	See "Deleting an MSDP storage server" on page 346.
Step 5	Delete the deduplication host configuration file	<p>Each load balancing server contains a deduplication host configuration file. If you use load balancing servers, delete the deduplication host configuration file from those servers.</p> <p>See "Deleting an MSDP host configuration file" on page 232.</p>
Step 6	Delete the credentials on deduplication servers	<p>If you have load balancing servers, delete the NetBackup Deduplication Engine credentials on those media servers.</p> <p>See "Deleting credentials from a load balancing server" on page 349.</p>

Table 11-1 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 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. ■ If you had configured the previous MSDP storage server to use MSDP Encryption using KMS service, you must use the same configuration for the new MSDP storage server. <p>See "About MSDP storage servers" on page 27.</p> <p>See "About MSDP server requirements" on page 28.</p>
Step 8	Connect the storage to the host	<p>Use the storage path that you configured for this replacement host.</p> <p>See the computer or the storage vendor's documentation.</p>
Step 9	Install the NetBackup media server software on the new host	<p>See the <i>NetBackup Installation Guide for UNIX and Windows</i>:</p> <p>http://www.veritas.com/docs/DOC5332</p>
Step 10	Reconfigure deduplication	<p>See "Configuring MSDP server-side deduplication" on page 55.</p>
Step 11	Import the backup images	<p>See the <i>NetBackup Administrator's Guide, Volume I</i>:</p> <p>http://www.veritas.com/docs/DOC5332</p>

Uninstalling MSDP

This chapter includes the following topics:

- [About uninstalling MSDP](#)
- [Deactivating MSDP](#)

About uninstalling MSDP

You cannot uninstall media server deduplication components separately from NetBackup. The deduplication components are installed when you install NetBackup software, and they are uninstalled when you uninstall NetBackup software.

Other topics describe related procedures, as follow:

- Reconfigure an existing deduplication environment.
See [“Changing the MSDP storage server name or storage path”](#) on page 343.
- Deactivate deduplication and remove the configuration files and the storage files.
See [“Deactivating MSDP”](#) on page 388.

Deactivating MSDP

You cannot remove the deduplication components from a NetBackup media server. You can disable the components and remove the deduplication storage files and the catalog files. The host remains a NetBackup media server.

This process assumes that all backup images that reside on the deduplication disk storage have expired.

Warning: If you remove deduplication and valid NetBackup images reside on the deduplication storage, data loss may occur.

Table 12-1 Remove MSDP

Step	Task	Procedure
Step 1	Remove client deduplication	Remove the clients that deduplicate their own data from the client deduplication list. See “Disabling MSDP client-side deduplication for a client” on page 127.
Step 2	Delete the storage units that use the disk pool	See the <i>NetBackup Administrator's Guide, Volume I</i> : http://www.veritas.com/docs/DOC5332
Step 3	Delete the disk pool	See “Deleting a Media Server Deduplication Pool” on page 361.
Step 4	Delete the deduplication storage server	See “Deleting an MSDP storage server” on page 346. 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 347.
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 232.
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 13-1 is a diagram of the storage server components.

Figure 13-1 MSDP server components




Table 13-1 describes the MSDP server components.

Table 13-1 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). ■ If required, compresses the data for backups and decompresses the backups for restores. ■ If required, encrypts the data for backups and decrypts the backups for restores. ■ If required, compresses the data for duplication and replication transfer. ■ If required, encrypts the data for duplication and replication transfer. <p>The plug-in runs on the deduplication storage server and on load balancing servers.</p>
Multi-Threaded Agent	<p>The NetBackup Deduplication Multi-Threaded Agent uses multiple threads for asynchronous network I/O and CPU core calculations. The agent runs on the storage server, load balancing servers, and clients that deduplicate their own data.</p> <p>See "About the MSDP Deduplication Multi-Threaded Agent" on page 59.</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 13-2.</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>nbostrpxy</code>) on the client.</p>
Reference database	<p>The reference database stores the references that point to every data segment of which a file is composed. Unique fingerprints identify data segments. The reference database is partitioned into multiple small reference database files to improve scalability and performance.</p> <p>The reference database is separate from the NetBackup catalog. The NetBackup catalog maintains the usual NetBackup backup image information.</p>

[Table 13-2](#) describes the components and functions within the NetBackup Deduplication Engine.

Table 13-2 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 363.</p>
Data Store Manager	<p>The Data Store Manager manages all of the data container files. The datastore Manager is a set of functions and threads that provides the following:</p> <ul style="list-style-type: none">■ A transaction mechanism to back up data into the datastore.■ A mechanism to read data from the datastore.■ A transaction mechanism to reclaim space in the datastore (that is, compact containers and remove containers). Container IDs are unique. The Data Store Manager increments the container number with each new container created. The data in a container is never overwritten, and a container ID is never reused. <p>See “About MSDP container files” on page 332.</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 68.</p>
Queue processing	<p>The NetBackup Deduplication Engine processes the transaction queue.</p> <p>See “About MSDP queue processing” on page 362.</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 13-2 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 13-2](#) diagram shows the backup process when a media server deduplicates the backups. The destination is a **Media Server Deduplication Pool**. A description follows.

Figure 13-2 Media server deduplication process



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

- The NetBackup Job Manager (`nbjm`) starts the Backup/Restore Manager (`bpbm`) on a media server.
- The Backup/Restore Manager starts the `bptm` process on the media server and the `bpbkar` process on the client.
- The Backup/Archive Manager (`bpbkar`) on the client generates the backup images and moves them to the media server `bptm` process.

The Backup/Archive Manager also sends the information about files within the image to the Backup/Restore Manager (`bpbrm`). The Backup/Restore Manager sends the file information to the `bpdbm` process on the master server for the NetBackup database.


- The `bptm` process moves the data to the deduplication plug-in.
 - The deduplication plug-in retrieves a list of IDs of the container files from the NetBackup Deduplication Engine. Those container files contain the fingerprints from the last full backup for the client. The list is used as a cache so the plug-in does not have to request each fingerprint from the engine.
 - The deduplication plug-in separates the files in the backup image into segments.
 - The deduplication plug-in buffers the segments and then sends batches of them to the Deduplication Multi-Threaded Agent. Multiple threads and shared memory are used for the data transfer.
 - The NetBackup Deduplication Multi-Threaded Agent processes the data segments in parallel using multiple threads to improve throughput performance. The agent then sends only the unique data segments to the NetBackup Deduplication Engine.
- If the host is a load balancing server, the Deduplication Engine is on a different host, the storage server.
- The NetBackup Deduplication Engine writes the data to the **Media Server Deduplication Pool**.

The first backup may have a 0% deduplication rate, although a 0% rate is unlikely. Zero percent means that all file segments in the backup data are unique.

The [Figure 13-3](#) diagram shows the backup process when a media server deduplicates the backups. The destination is a **PureDisk Deduplication Pool**. A description follows.

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

Figure 13-3 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 13-3 describes the client deduplication components.

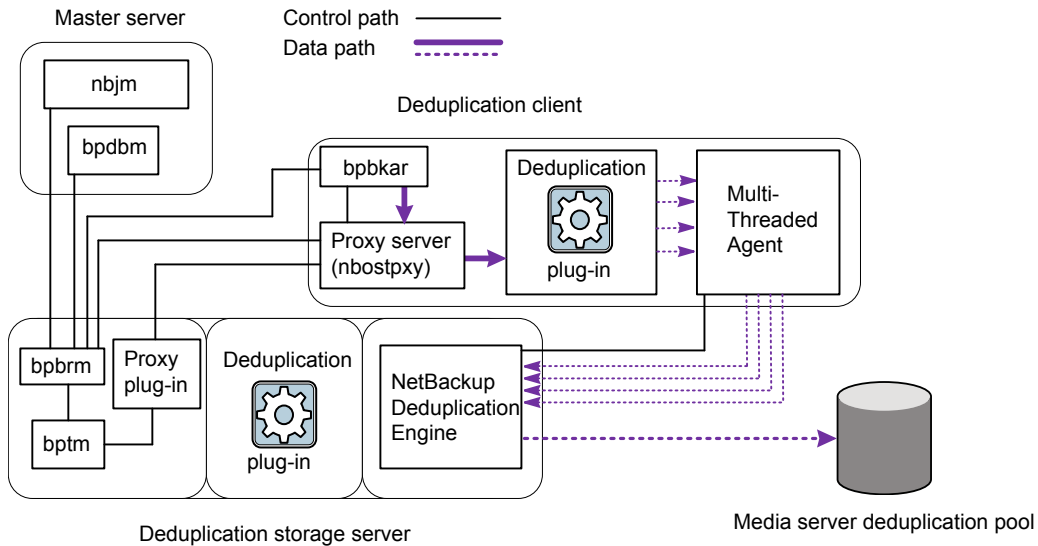
Table 13-3 Client MSDP components

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).■ If required, compresses the data for backups and decompresses the backups for restores.■ If required, encrypts the data for backups and decrypts the backups for restores.
Multi-Threaded Agent	The NetBackup Deduplication Multi-Threaded Agent uses multiple threads for asynchronous network I/O and CPU core calculations. The agent runs on the storage server, load balancing servers, and clients that deduplication their own data. See “About the MSDP Deduplication Multi-Threaded Agent” on page 59.
Proxy server	The OpenStorage proxy server (<code>nboostpxy</code>) manages control communication with the proxy plug-in on the storage server.

MSDP client–side deduplication backup process

The [Figure 13-4](#) 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 13-4 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 70.
 - 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 70.

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 13-5](#) diagram shows the backup process of a client that deduplicates its own data. The destination is a PureDisk storage pool. A description follows.

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

Figure 13-5 MSDP deduplication client backup to a PureDisk storage pool



The following list describes the backup process for an MSDP client to a PureDisk storage pool:

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

- The deduplication plug-in retrieves a list of fingerprints from the last full backup of the client. The list is used as a cache so the plug-in does not have to request each fingerprint from the engine.
- The deduplication plug-in performs file fingerprint calculations.
- The deduplication plug-in sends only unique data segments to the PureDisk storage pool.

Troubleshooting

This chapter includes the following topics:

- [About unified logging](#)
- [About legacy logging](#)
- [NetBackup 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](#)
- [Troubleshooting Cloud Catalyst issues](#)
- [Unable to obtain the administrator password to use an AWS EC2 instance that has a Windows OS](#)
- [Trouble shooting multi-domain issues](#)

About unified logging

Unified logging creates log file names and messages in a format that is standardized across Veritas products. Only the `vxlogview` command can assemble and display the log information correctly. Server processes and client processes use unified logging.

Log files for originator IDs are written to a subdirectory with the name specified in the log configuration file. All unified logs are written to subdirectories in the following directory:

Windows `install_path\NetBackup\logs`

UNIX `/usr/opensv/logs`

You can access logging controls in **Logging** host properties. You can also manage unified logging with 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 404.
for more information about the `vxlogview` command.

About using the `vxlogview` command to view unified logs

Only the `vxlogview` command can assemble and display the unified logging information correctly. The unified logging files are in binary format and some of the information is contained in an associated resource file. These logs are stored in the following directory. You can display `vxlogview` results faster by restricting the search to the files of a specific process.

UNIX `/usr/opensv/logs`

Windows `install_path\NetBackup\logs`

Table 14-1 Fields in `vxlogview` query strings

Field name	Type	Description	Example
PRODID	Integer or string	Provide the product ID or the abbreviated name of product.	PRODID = 51216 PRODID = 'NBU'
ORGID	Integer or string	Provide the originator ID or the abbreviated name of the component.	ORGID = 116 ORGID = 'nbpem'
PID	Long Integer	Provide the process ID	PID = 1234567

Table 14-1 Fields in vxlogview query strings (*continued*)

Field name	Type	Description	Example
TID	Long Integer	Provide the thread ID	TID = 2874950
STDATE	Long Integer or string	Provide the start date in seconds or in the locale-specific short date and time format. For example, a locale can have the format 'mm/dd/yy hh:mm:ss AM/PM'	STDATE = 98736352 STDATE = '4/26/11 11:01:00 AM'
ENDATE	Long Integer or string	Provide the end date in seconds or in the locale-specific short date and time format. For example, a locale can have the format 'mm/dd/yy hh:mm:ss AM/PM'	ENDATE = 99736352 ENDATE = '04/27/11 10:01:00 AM'
PREVTIME	String	Provide the hours in 'hh:mm:ss' format. This field should be used only with operators =, <, >, >=, and <=	PREVTIME = '2:34:00'
SEV	Integer	Provide one of the following possible severity types: 0 = INFO 1 = WARNING 2 = ERR 3 = CRIT 4 = EMERG	SEV = 0 SEV = INFO
MSGTYPE	Integer	Provide one of the following possible message types: 0 = DEBUG (debug messages) 1 = DIAG (diagnostic messages) 2 = APP (application messages) 3 = CTX (context messages) 4 = AUDIT (audit messages)	MSGTYPE = 1 MSGTYPE = DIAG

Table 14-1 Fields in vxlogview query strings (*continued*)

Field name	Type	Description	Example
CTX	Integer or string	Provide the context token as string identifier or 'ALL' to get all the context instances to be displayed. This field should be used only with the operators = and !=.	CTX = 78 CTX = 'ALL'

Table 14-2 Examples of query strings with dates

Example	Description
(PRODID == 51216) && ((PID == 178964) ((STDATE == '2/5/15 09:00:00 AM') && (ENDATE == '2/5/15 12:00:00 PM')))	Retrieves the log file message for the NetBackup product ID 51216 between 9AM and 12PM on 2015-05-02.
((prodid = 'NBU') && ((stdate >= '11/18/14 00:00:00 AM') && (enddate <= '12/13/14 12:00:00 PM')) ((prodid = 'BENT') && ((stdate >= '12/12/14 00:00:00 AM') && (enddate <= '12/25/14 12:00:00 PM')))	Retrieves the log messages for the NetBackup product NBU between 2014-18-11 and 2014-13-12 and the log messages for the NetBackup product BENT between 2014-12-12 and 2014-25-12.
(STDATE <= '04/05/15 0:0:0 AM')	Retrieves the log messages that were logged on or before 2015-05-04 for all of the installed Veritas products.

Examples of using vxlogview to view unified logs

The following examples demonstrate how to use the vxlogview command to view unified logs.

Table 14-3 Example uses of the vxlogview command

Item	Example
Display all the attributes of the log messages	vxlogview -p 51216 -d all

Table 14-3 Example uses of the vxlogview command (*continued*)

Item	Example
Display specific attributes of the log messages	<p>Display the log messages for NetBackup (51216) that show only the date, time, message type, and message text:</p> <pre>vxlogview --prodid 51216 --display D,T,m,x</pre>
Display the latest log messages	<p>Display the log messages for originator 116 (nbpem) that were issued during the last 20 minutes. Note that you can specify <code>-o nbpem</code> instead of <code>-o 116</code>:</p> <pre># vxlogview -o 116 -t 00:20:00</pre>
Display the log messages from a specific time period	<p>Display the log messages for nbpem that were issued during the specified time period:</p> <pre># vxlogview -o nbpem -b "05/03/15 06:51:48 AM" -e "05/03/15 06:52:48 AM"</pre>
Display results faster	<p>You can use the <code>-i</code> option to specify an originator for a process:</p> <pre># vxlogview -i nbpem</pre> <p>The <code>vxlogview -i</code> option searches only the log files that the specified process (nbpem) 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>
Search for a job ID	<p>You can search the logs for a particular job ID:</p> <pre># vxlogview -i nbpem grep "jobid=job_ID"</pre> <p>The <code>jobid=</code> search key should contain no spaces and must be lowercase.</p> <p>When searching for a job ID, you can use any <code>vxlogview</code> command option. This example uses the <code>-i</code> option with the name of the process (nbpem). 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>

About legacy logging

In NetBackup legacy debug logging, a process creates log files of debug activity in its own logging directory. By default, NetBackup creates only a subset of logging directories, in the following locations:

Windows	<code>install_path\NetBackup\logs</code> <code>install_path\Volmgr\debug</code>
UNIX	<code>/usr/opensv/netbackup/logs</code> <code>/usr/opensv/volmgr/debug</code>

It is recommended that you do not use symbolic links or hard links inside legacy log folders.

If any process runs for a non-root or non-admin user and there is no logging happening under legacy log folders, you can create a folder using the `mklogdir` command for the required user.

To run a command line for a non-root or non-admin user (troubleshooting when the NetBackup services are not running), it is recommended that you create user folders for the specific command line. You can create the folders either using the `mklogdir` command or manually with the non-root or non-admin user privileges.

To use legacy logging, a log file directory must exist for a process. If the directory is not created by default, you can use the Logging Assistant or the `mklogdir` batch files to create the directories. Or, you can manually create the directories. When logging is enabled for a process, a log file is created when the process begins. Each log file grows to a certain size before the NetBackup process closes it and creates a new log file.

You can use the following batch files to create all of the log directories:

- Windows: `install_path\NetBackup\Logs\mklogdir.bat`
- UNIX: `/usr/opensv/netbackup/logs/mklogdir`

More information

See the [NetBackup Commands Reference Guide](#) for a complete description about the `mklogdir` command.

Creating NetBackup log file directories for MSDP

Before you configure your NetBackup feature, create the directories into which the NetBackup commands write log files. Create the directories on the master server

and on each media server that you use for your feature. The log files reside in the following directories:

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

More information about NetBackup logging is available in the *NetBackup Logging Reference Guide*, available through the following URL:

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

To create log directories for NetBackup commands

- ◆ Depending on the operating system, run one of the following scripts:

UNIX: `/usr/opensv/netbackup/logs/mklogdir`

Windows: `install_path\NetBackup\logs\mklogdir.bat`

To create the `tpconfig` command log directory

- ◆ Depending on the operating system, create the `debug` directory and the `tpcommand` directory (by default, the `debug` directory and the `tpcommand` directory do not exist). The pathnames of the directories are as follows:

UNIX: `/usr/opensv/volmgr/debug/tpcommand`

Windows: `install_path\Veritas\Volmgr\debug\tpcommand`

NetBackup 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 Veritas Unified Logging (VxUL) log 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 406.

See “[About unified logging](#)” on page 401.

In VxUL logs, the messages that begin with an `sts` prefix relate to the interaction with the deduplication plug-in. Most interaction occurs on the NetBackup media servers. To view and manage VxUL log files, you must use NetBackup log commands. For information about how to use and manage logs on NetBackup servers, see the *NetBackup Logging Reference Guide*. The guide is available through the following URL:

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

Most interaction occurs on the NetBackup media servers. Therefore, the log files on the media servers that you use for disk operations are of most interest.

Warning: The higher the log level, the greater the affect on NetBackup performance. Use a log level of 5 (the highest) only when directed to do so by a Veritas 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 14-4](#).

[Table 14-4](#) describes the log files for each component.

Table 14-4 Logs for NetBackup MSDP activity

Component	VxUL OID	Description
Backups and restores	117	The nbjm Job Manager.
Backups and restores	N/A	<p>Messages appear in the log files for the following processes:</p> <ul style="list-style-type: none"> ■ The bpbrm backup and restore manager. The following is the path to the log files: UNIX: <code>/usr/opensv/netbackup/logs/bpbrm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bpbrm</code> ■ The bpdbrm database manager. The following is the path to the log files: UNIX: <code>/usr/opensv/netbackup/logs/bpdbrm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bpdbrm</code> ■ The bptm tape manager for I/O operations. The following is the path to the log files: UNIX: <code>/usr/opensv/netbackup/logs/bptm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bptm</code>

Table 14-4 Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
Catalog shadow copies	N/A	<p>The MSDP catalog shadow copy process writes messages to the following log files and directories:</p> <p>UNIX:</p> <pre>/storage_path/log/spad/spad.log /storage_path/log/spad/sched_CatalogBackup.log /storage_path/log/spad/client_name/</pre> <p>Windows:</p> <pre>storage_path\log\spad\spad.log storage_path\log\spad\sched_CatalogBackup.log storage_path\log\spad\client_name\</pre>
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>: <ul style="list-style-type: none"> UNIX: <code>/usr/opensv/netbackup/logs/bptm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bptm</code> ■ For <code>bpstsinfo</code>: <ul style="list-style-type: none"> Windows: <code>/usr/opensv/netbackup/logs/admin</code> UNIX: <code>/usr/opensv/netbackup/logs/bpstsinfo</code> Windows: <code>install_path\Veritas\NetBackup\logs\admin</code> Windows: <code>install_path\Veritas\NetBackup\logs\stsinfo</code> ■ For <code>bpbrm</code>: <ul style="list-style-type: none"> UNIX: <code>/usr/opensv/netbackup/logs/bpbrm</code> Windows: <code>install_path\Veritas\NetBackup\logs\bpbrm</code>
Client deduplication proxy server	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>

Table 14-4 Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
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>The <code>DEBUGLOG</code> entry and the <code>LOGLEVEL</code> in the <code>pd.conf</code> file determine the log location and level for the deduplication plug-in. The following are the default locations for log files:</p> <ul style="list-style-type: none"> ■ UNIX: <code>/var/log/puredisk/pdplugin.log</code> ■ Windows: <code>C:\pdplugin.log</code> <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> entry in the <code>pd.conf</code> file.</p> <p>See “About the MSDP pd.conf configuration file” on page 209.</p> <p>See “Editing the MSDP pd.conf file” on page 210.</p>
Device configuration and monitoring	111	The <code>nbemm</code> process.
Device configuration and monitoring	178	The Disk Service Manager process that runs in the Enterprise Media Manager (EMM) process.
Device configuration and monitoring	202	The storage server interface process that runs in the Remote Manager and Monitor Service. RMMS runs on media servers.
Device configuration and monitoring	230	The Remote Disk Service Manager interface (RDSM) that runs in the Remote Manager and Monitor Service. RMMS runs on media servers.

Table 14-4 Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
drcontrol utility	N/A	<p>You must run the <code>drcontrol</code> utility on the MSDP storage server host. The command requires administrator privileges.</p> <p>The utility creates a log file and displays its pathname in the command output. The utility writes log files to the following directory, depending on the operating system:</p> <p>UNIX:</p> <pre>/[storage_path]/log/drcontrol/policy_admin</pre> <pre>/storage_path/log/drcontrol/dedupe_catalog_DR</pre> <p>Windows:</p> <pre>storage_path\log\drcontrol\policy_admin</pre> <pre>storage_path\log\drcontrol\dedupe_catalog_DR</pre> <p>See “About protecting the MSDP catalog” on page 233.</p> <p>See “About recovering the MSDP catalog” on page 377.</p>
Installation	N/A	<p>The NetBackup installation process writes information about the installation of the deduplication components to a log file in the following directory:</p> <ul style="list-style-type: none"> ■ UNIX: <code>/var/log/puredisk</code> ■ Windows: <code>%ALLUSERSPROFILE%\Symantec\NetBackup\InstallLogs</code>
NetBackup Deduplication Engine	N/A	<p>The NetBackup Deduplication Engine writes several log files, as follows:</p> <ul style="list-style-type: none"> ■ Log files in the <code>storage_path/log/spoold</code> directory, as follows: <ul style="list-style-type: none"> ■ The <code>spoold.log</code> file is the main log file ■ The <code>storaged.log</code> file is for queue processing. ■ A log file for each connection to the engine is stored in a directory in the storage path <code>spoold</code> directory. The following describes the pathname to a log file for a connection: <pre>hostname/application/TaskName/MMDDYY.log</pre> <p>For example, the following is an example of a <code>crcontrol</code> connection log pathname on a Linux system:</p> <pre>/storage_path/log/spoold/server.example.com/crcontrol/Control/010112.log</pre> <p>Usually, the only reason to examine these connection log files is if a Veritas support representative asks you to.</p> ■ A VxUL log file for the events and errors that NetBackup receives from polling. The originator ID for the deduplication engine is 364.

Table 14-4 Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
NetBackup Deduplication Engine	364	The NetBackup Deduplication Engine that runs on the deduplication storage server.
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: <i>hostname/application/TaskName/MMDDYY.log</i> For example, the following is an example of a <code>bpstsinfo</code> connection log pathname on a Linux system: <i>/storage_path/log/spoold/server.example.com/bpstsinfo/spad/010112.log</i> Usually, the only reason to examine these connection log files is if a Veritas support representative asks you to. <p>You can set the log level and retention period in the Change Storage Server dialog box Properties tab.</p> <p>See “Changing MSDP storage server properties” on page 341.</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: <i>/storage_path/log/spad/replication.log</i>

Table 14-4 Logs for NetBackup MSDP activity (*continued*)

Component	VxUL OID	Description
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, Veritas recommends that you set the logging level for OID 387 to 2 or less. To configure unified logs, see the following guide:</p> <p>The <i>NetBackup Logging Reference Guide</i>:</p> <p>http://www.veritas.com/docs/DOC5332</p>
Resilient network connections	N/A	<p>The deduplication plug-in logs information about keeping the connection alive. For more information about the deduplication plug-in log file, see “Deduplication plug-in” in this table.</p> <p>The <code>pd.conf</code> file <code>FILE_KEEP_ALIVE_INTERVAL</code> parameter controls the connection keep alive interval.</p> <p>See “About the MSDP pd.conf configuration file” on page 209.</p> <p>See “Editing the MSDP pd.conf file” on page 210.</p>

Troubleshooting MSDP installation issues

The following sections may help you troubleshoot installation issues.

See “[MSDP installation on SUSE Linux fails](#)” on page 413.

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 [“NetBackup MSDP log files”](#) on page 407.

See [“MSDP storage server configuration fails”](#) on page 414.

See [“MSDP database system error \(220\)”](#) on page 414.

See [“MSDP server not found error”](#) on page 415.

See [“License information failure during MSDP configuration”](#) on page 416.

See [“The disk pool wizard does not display an MSDP volume”](#) on page 416.

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 232.

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	<p>RDSM has encountered an STS error:</p> <pre>Failed to update storage server <i>ssname</i>, database system error</pre>
Diagnosis	<p>The <code>PDDE_initConfig</code> script was invoked, but errors occurred during the storage initialization.</p> <p>First, examine the deduplication configuration script log file for references to the server name.</p> <p>See “NetBackup MSDP log files” on page 407.</p> <p>Second, examine the <code>tpconfig</code> command log file errors about creating the credentials for the server name. The <code>tpconfig</code> command writes to the standard NetBackup administrative commands log directory.</p>

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	<p>RDSM has encountered an issue with STS where the server was not found: <code>getStorageServerInfo</code></p> <pre>Failed to create storage server <i>ssname</i>, invalid command parameter</pre>
Diagnosis	<p>Possible root causes:</p> <ul style="list-style-type: none"> ■ When you configured the storage server, you selected a media server that runs an unsupported operating system. All media servers in your environment appear in the Storage Server Configuration Wizard; be sure to select only a media server that runs a supported operating system. ■ If you used the <code>nbdevconfig</code> command to configure the storage server, you may have typed the host name incorrectly. Also, case matters for the storage server type, so ensure that you use PureDisk for the storage server type.

License information failure during 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 417.

See [“MSDP backup or duplication job fails”](#) on page 417.

See [“MSDP client deduplication fails”](#) on page 420.

See [“MSDP volume state changes to DOWN when volume is unmounted”](#) on page 420.

See [“MSDP errors, delayed response, hangs”](#) on page 421.

See [“Cannot delete an MSDP disk pool”](#) on page 422.

See [“MSDP media open error \(83\)”](#) on page 423.

See [“MSDP media write error \(84\)”](#) on page 425.

See [“MSDP no images successfully processed \(191\)”](#) on page 428.

See [“MSDP storage full conditions”](#) on page 428.

See [“Troubleshooting MSDP catalog backup”](#) on page 429.

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 28.

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 430.

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 362.

See [“NetBackup MSDP log files”](#) on page 407.

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 [“NetBackup MSDP log files”](#) on page 407.

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, Veritas recommends that you use fully-qualified domain names for your NetBackup environment.

See [“Use fully qualified domain names”](#) on page 41.

MSDP volume state changes to DOWN when volume is unmounted

If a volume becomes unmounted, NetBackup changes the volume state to DOWN. NetBackup jobs that require that volume fail.

To determine the volume state

- ◆ Invoke the following command on the master server or the media server that functions as the deduplication storage server:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbdevquery -listdv -stype PureDisk -U`

Windows: `install_path\NetBackup\bin\admincmd\nbdevquery -listdv -stype PureDisk -U`

The following example output shows that the `DiskPoolVolume` is UP:

```
Disk Pool Name      : PD_Disk_Pool
Disk Type           : PureDisk
Disk Volume Name    : PureDiskVolume
Disk Media ID       : @aaaaab
Total Capacity (GB) : 49.98
Free Space (GB)     : 43.66
Use%                : 12
Status              : UP
Flag                : ReadOnWrite
Flag                : AdminUp
Flag                : InternalUp
Num Read Mounts     : 0
Num Write Mounts    : 1
Cur Read Streams   : 0
Cur Write Streams  : 0
Num Repl Sources    : 0
Num Repl Targets    : 0
WORM Lock Min Time  : 0
WORM Lock Max Time  : 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 359.

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 28.

For virtual machines, Veritas 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, Veritas recommends that you try this solution first when troubleshooting performance issues.

Cannot delete an MSDP disk pool

If you cannot delete a disk pool that you believe contains no valid backup images, the following information may help you troubleshoot the problem.

- [Expired fragments remain on MSDP disk](#)
- [Incomplete SLP duplication jobs](#)

Expired fragments remain on MSDP disk

Under some circumstances, the fragments that compose an expired backup image may remain on disk even though the images have expired. For example, if the storage server crashes, normal clean-up processes may not run. In those circumstances, you cannot delete a disk pool because image fragment records still exist. The error message may be similar to the following:

```
DSM has found that one or more volumes in the disk pool diskpoolname
has image fragments.
```

To delete the disk pool, you must first delete the image fragments. The `nbdelete` command deletes expired image fragments from disk volumes.

To delete the fragments of expired images

- ◆ Run the following command on the master server:

```
UNIX: /usr/opensv/netbackup/bin/admincmd/nbdelete -allvolumes -force
```

```
Windows: install_path\NetBackup\bin\admincmd\nbdelete -allvolumes
-force
```

The `-allvolumes` option deletes expired image fragments from all volumes that contain them.

The `-force` option removes the database entries of the image fragments even if fragment deletion fails.

Incomplete SLP duplication jobs

Incomplete storage lifecycle policy duplication jobs may prevent disk pool deletion. You can determine if incomplete jobs exist and then cancel them.

To cancel storage lifecycle policy duplication jobs

- 1 Determine if incomplete SLP duplication jobs exist by running the following command on the master server:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbstlutil stlilist -image_incomplete`

Windows: `install_path\NetBackup\bin\admincmd\nbstlutil stlilist -image_incomplete`

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

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

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

MSDP media open error (83)

The `media open error (83)` message is a generic error for the duplication. The error appears in the **NetBackup Administration Console Activity Monitor**.

Often, the NetBackup Deduplication Engine (`spoold`) or the NetBackup Deduplication Manager (`spad`) were too busy to respond to the deduplication process in a timely manner. External factors may cause the Deduplication Engine or the Deduplication Manager to be unresponsive. Were they temporarily busy (such as queue processing in progress)? Do too many jobs run concurrently?

See [“About MSDP performance”](#) on page 38.

Usually but not always the NetBackup `bpdm` log provides additional information about status 83.

The following subsections describe use cases that generated an error 83.

SQL Server client-side backups fail

Client-side backups of a SQL Server database may fail in the following circumstances:

- The **Both IPv4 and IPv6** option is enabled for the master server, the media server that hosts the NetBackup Deduplication Engine, and the client. The **Both IPv4 and IPv6** option is configured in the **Network Settings** host properties.
- The IPv6 network is configured as a preferred network for the master server, the media server that hosts the NetBackup Deduplication Engine, and the client. The preferred network **Match (Above network will be preferred for communication)** property also is enabled. Preferred networks are configured in the **Preferred Networks** host properties.
- The IPv6 network is chosen for the backup.

Examine the `bpbrm` log file for an error similar to the following:

```
probe_ost_plugin: sts_get_server_prop_byname failed: error 2060057
```

If the error message appears, the NetBackup host name cache may not contain the correct host name mapping information. The cache may be out of sync if DNS changes in your network environment were not fully propagated throughout your environment. It takes some amount of time for DNS changes to propagate throughout a network environment.

To resolve the problem, do the following on the NetBackup master server and on the MSDP storage server:

1. Stop the NetBackup services.
2. Run the following command:

UNIX: `/usr/opensv/netbackup/bin/bpclntcmd -clearhostcache`

Windows: `install_path\NetBackup\bin\bpclntcmd.exe -clearhostcache`

3. Start the NetBackup services.

For more information about client deduplication logging, see the description of “Client deduplication proxy plug-in” in the “MSDP log files” topic.

See [“NetBackup MSDP log files”](#) on page 407.

Restore or duplication fails

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

[Table 14-5](#) describes other messages that may appear.

Table 14-5 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 (NetBackup 5000 series appliance source)	Image open failed: error 2060018: file not found	<code>sts_open_image</code> failed: error 2060018
Replication (NetBackup 5000 series appliance source)	get image properties failed: error 2060013: no more entries	<code>rpl_add_image_set:</code> <code>rpl_get_image_info()</code> failed, error 2060013

The messages may indicate a client name case sensitivity issue in your MSDP environment. For a description of the problem and the procedures to resolve it, see the following Veritas tech note:

<http://www.veritas.com/docs/TECH207194>

MSDP media write error (84)

[Table 14-6](#) describes solutions to the media write errors that may occur during **Media Server Deduplication Pool** backups, duplication, or replication.

Also see the following subsections for descriptions of more complicated solutions:

- [Host name resolution problems](#)

Table 14-6 Media write error causes

The NetBackup Deduplication Engine (spoold) was too busy to respond.	Examine the Disk Logs report for errors that include the name PureDisk. Examine the disk monitoring services log files for details from the deduplication plug-in. See “Viewing MSDP disk reports” on page 334.
Data removal is running.	Data cannot be backed up at the same time as it is removed. See “About MSDP queue processing” on page 362.
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 49.
The deduplication pool is down.	Change the state to up. See “Changing Media Server Deduplication Pool state” on page 351.
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 231.

To fix this problem, configure your network environment so that all permutations of the storage server name resolve.

Veritas recommends that you use fully qualified domain names.

See [“Use fully qualified domain names”](#) on page 41.

Duplication job to PureDisk storage fails

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

A duplication job to PureDisk storage may fail for a large number of small files. This failure may occur if the storage unit maximum fragment is at the default setting.

Messages similar to the following may appear in the NetBackup job details:

```
2/12/2013 2:33:20 AM - Critical bpdm(pid=25204) Storage Server Error:
  (Storage server: PureDisk: msdp.example.com) 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.example.com backup id msdp.example.com_1360303271 optimized
  duplication failed, media write error (84).
2/12/2013 2:33:39 AM - Info msdp.example.com(pid=25204)
  StorageServer=PureDisk:msdp.example.com; Report=PDDO Stats for
  (msdp.example.com): 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.example.com_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 14-7](#) describes other messages that may appear.

Table 14-7 Case sensitivity error messages

Operation	Activity Monitor job details	Status in <code>bpdm</code> and <code>bptm</code> log files
Verify	image open failed: error 2060018: file not found	sts_open_image failed: error 2060018

The message may indicate a client name case sensitivity issue in your MSDP environment. For a description of the problem and the procedures to resolve it, see the following Veritas tech note:

<http://www.veritas.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.

Note: In a Cloud Catalyst environment, however, the UNIX `df` command can be used to determine the disk space in the local cache directory.

See [“About MSDP storage capacity and usage reporting”](#) on page 330.

See [“About MSDP container files”](#) on page 332.

See [“Viewing storage usage within MSDP container files”](#) on page 332.

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 362.

See “[About the MSDP data removal process](#)” on page 371.

Although not advised, you can reclaim free space manually.

See “[Processing the MSDP transaction queue manually](#)” on page 362.

Troubleshooting MSDP catalog backup

The following subsections provide information about MSDP catalog backup and recovery.

Catalog backup

[Table 14-8](#) describes error messages that may occur when you create or update a catalog backup policy. The messages are displayed in the shell window in which you ran the `drcontrol` utility. The utility also writes the messages to its log file.

Table 14-8 MSDP `drcontrol` codes and messages

Code or message	Description
1	Fatal error in an operating system or deduplication command that the <code>drcontrol</code> utility calls.
110	The command cannot find the necessary NetBackup configuration information.
140	The user who invoked the command does not have administrator privileges.
144	A command option or argument is required.
226	The policy name that you specified already exists.
227	This error code is passed from the NetBackup <code>bplist</code> command. The MSDP catalog backup policy you specified does not exist or no backups exist for the given policy name.
255	Fatal error in the <code>drcontrol</code> utility.

For more information about status codes and error messages, see the following:

- The Troubleshooter in the NetBackup Administration Console.
- The *NetBackup Status Codes Reference Guide* available through the following webpage:
<http://www.veritas.com/docs/DOC5332>

Catalog recovery from a shadow copy

If NetBackup detects corruption in the MSDP catalog, the Deduplication Manager recovers the catalog automatically from the most recent shadow copy. That recovery process also plays a transaction log so that the recovered MSDP catalog is current.

Although the shadow copy recovery process is automatic, a recovery procedure is available if you need to recover from a shadow copy manually.

See [“Restoring the MSDP catalog from a shadow copy”](#) on page 378.

Storage Platform Web Service (spws) does not start

Storage Platform Web Service (spws) does not start when you run `bp.start_all`.

Workaround:

If spws does not start when you run `bp.start_all`, run the following command to reconfigure vpfs and spws:

```
vpfs_config.sh --configure_byo
```

Disk volume API or command line option does not work

You have an MSDP storage server that has NetBackup version earlier than 8.3 and you have not enabled the encryption and KMS details. If you try to update the encryption and KMS details for a local volume using the new disk volume update API, the API operation succeeds. However, the actual values are not updated.

This issue occurs for both the API and command line option.

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 334.
- The NetBackup `bpperor` 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 14-9 MSDP event codes and messages

Event #	Event Severity	NetBackup Severity	Message example
1000	2	Error	Operation configload/reload failed on server PureDisk:server1.example.com on host server1.example.com.
1001	2	Error	Operation configload/reload failed on server PureDisk:server1.example.com on host server1.example.com.
1002	4	Warning	The open file limit exceeded in server PureDisk:server1.example.com on host server1.example.com. Will attempt to continue further.
1003	2	Error	A connection request was denied on the server PureDisk:server1.example.com on host server1.example.com.
1004	1	Critical	Network failure occurred in server PureDisk:server1.example.com on host server1.example.com.
1008	2	Error	Task Aborted; An unexpected error occurred during communication with remote system in server PureDisk:server1.example.com on host server1.example.com.
1009	8	Authorization	Authorization request from <IP> for user <USER> denied (<REASON>).
1010	2	Error	Task initialization on server PureDisk:server1.example.com on host server1.example.com got an unexpected error.
1011	16	Info	Task ended on server PureDisk:server1.example.com on host server1.example.com.
1013	1	Critical	Task session start request on server PureDisk:server1.example.com on host server1.example.com got an unexpected error.
1012	2	Error	A request for agent task was denied on server PureDisk:server1.example.com on host server1.example.com.

Table 14-9 MSDP event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1014	1	Critical	Task session start request on server PureDisk:server1.example.com on host server1.example.com got an unexpected error.
1015	1	Critical	Task creation failed, could not initialize task class on server PureDisk:server1.example.com on host server1.example.com.
1017	1	Critical	Service Veritas DeduplicationEngine exit on server PureDisk:server1.example.com on host server1.example.com. Please check the server log for the probable cause of this error. The application has terminated.
1018	16	Info	Startup of Veritas Deduplication Engine completed successfully on server1.example.com.
1019	1	Critical	Service Veritas DeduplicationEngine restart on server PureDisk:server1.example.com on host server1.example.com. Please check the server log for the probable cause of this error. The application has restarted.
1020	1	Critical	Service Veritas Deduplication Engine connection manager restart failed on server PureDisk:server1.example.com on host server1.example.com. Please check the server log for the probable cause of this error.The application has failed to restart.
1028	1	Critical	Service Veritas DeduplicationEngine abort on server PureDisk:server1.example.com on host server1.example.com. 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.example.com in server server1.example.com.
1030	1	Critical	Operation Storage Database Initialization failed on server PureDisk:server1.example.com on host server1.example.com.
1031	1	Critical	Operation Content router context initialization failed on server PureDisk:server1.example.com on host server1.example.com.

Table 14-9 MSDP event codes and messages (*continued*)

Event #	Event Severity	NetBackup Severity	Message example
1032	1	Critical	Operation log path creation/print failed on server PureDisk:server1.example.com on host server1.example.com.
1036	4	Warning	Operation a transaction failed on server PureDisk:server1.example.com on host server1.example.com.
1037	4	Warning	Transaction failed on server PureDisk:server1.example.com on host server1.example.com. Transaction will be retried.
1040	2	Error	Operation Database recovery failed on server PureDisk:server1.example.com on host server1.example.com.
1043	2	Error	Operation Storage recovery failed on server PureDisk:server1.example.com on host server1.example.com.
1044	multiple	multiple	The usage of one or more system resources has exceeded a warning level. Operations will or could be suspended. Please take action immediately to remedy this situation.
1057			A data corruption has been detected. The data consistency check detected a data loss or data corruption in the Media Server Deduplication Pool (MSDP) and reported the affected backups. The backup ID and policy name appear in the NetBackup Disk Logs report and the <code>storage_path/log/spoold/storaged.log</code> file on the storage server.
2000		Error	Low space threshold exceeded on the partition containing the storage database on server PureDisk:server1.example.com on host server1.example.com.

Troubleshooting Cloud Catalyst issues

The following sections may help you troubleshoot issues that involve NetBackup Cloud Catalyst.

See [“Cloud Catalyst logs”](#) on page 434.

See [“Problems encountered while using the Cloud Storage Server Configuration Wizard”](#) on page 435.

See [“Disk pool problems”](#) on page 437.

See [“Problems during cloud storage server configuration”](#) on page 438.

- See [“Status 191: No images were successfully processed”](#) on page 440.
- See [“Media write error \(84\) if due to a full local cache directory”](#) on page 440.
- See [“Restarting the vxesfsd process”](#) on page 441.
- See [“Problems restarting vxesfsd”](#) on page 443.
- See [“Cloud Catalyst troubleshooting tools”](#) on page 444.
- See [“Unable to create CloudCatalyst with a media server that has version earlier to 8.2”](#) on page 443.

Cloud Catalyst logs

In addition to the MSDP logs, NetBackup Cloud Catalyst creates additional logs.

The logs for the `vxesfsd` and the `ocsd` process are located in the following directories:

- `/usr/openv/netbackup/logs/esfs_database`
- `/usr/openv/netbackup/logs/esfs_filesystem`
- `/usr/openv/netbackup/logs/esfs_monitor`
- `/usr/openv/netbackup/logs/esfs_storage`
- `/usr/openv/netbackup/logs/esfs_storagemanager`

The location of the logs is configurable in the `esfs.json` file.

See [“About the Cloud Catalyst esfs.json configuration file”](#) on page 271.

Additional logs (start and stop operations, configuration, upgrade, and more) are found in:

- `/usr/openv/esfs/log/ops`

Error messages in `esfs_filesystem`

The following error message may be found in the `esfs_filesystem` log:

`/usr/openv/netbackup/logs/esfs_filesystem:`

Error message	ERROR - "Failed to write /data/521/534297.bhd allow_write F skip_write_fail F"
---------------	---

Diagnosis

One possible cause is that cache eviction cannot clean up any data from the local cache directory for non-MSDP system file writes.

The cache eviction is triggered when the high watermark setting (`HighWatermark`) for the directory is reached. Once the high watermark is reached, data is purged when the used space reaches the midpoint between `HighWatermark` and `LowWatermark` $(high+low)/2$ and continues until `LowWatermark` is reached. If the rate of incoming data exceeds the rate where the watermark can be maintained, the jobs begin to fail.

Since the upload speed to the cloud is slower than the hard disk drive write speed, this situation can arise when the partition for the local cache directory is small. Increase the space that is allowed for the local cache directory.

Problems encountered while using the Cloud Storage Server Configuration Wizard

The following topics are problems that a user might encounter while using the Cloud Storage Server Configuration Wizard to configure a Cloud Catalyst storage server.

Unable to select the desired media server

In the **Add Storage Server** wizard screen, the user selects a media server in the drop-down menu that will become the Cloud Catalyst storage server.

If the media server that you want to use does not appear in the drop-down menu, check the following:

Media server version

The NetBackup media servers that you use as Cloud Catalyst storage servers must be the same NetBackup version as the master server.

Host name-based security certificate

This is applicable for media server versions 8.0 to 8.1.2 only.

If the media server that you want to use does not appear in the drop-down menu, make sure that a host name-based security certificate has been deployed for that media server. Use the following procedure to deploy a certificate if needed.


1. `/usr/opensv/netbackup/bin/admincmd/bpnbaz -ProvisionCert host_name`
2. Restart the services on the media server after generating a certificate.

Unable to select the Deduplication option

To create a Cloud Catalyst storage server, the **Enable NetBackup Cloud Catalyst** option must be selected in the Cloud Storage Server Configuration Wizard. Also, a directory must be indicated as the **Local cache directory**.

These options are grayed out if any of the following cases are true:

- The selected media server does not have NetBackup 8.1 or later installed.
- Cloud Catalyst does not support the media server operating system.
- Cloud Catalyst does not support the cloud vendor.



See the NetBackup compatibility lists for support information:


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

Storage Server Creation Status errors

The following errors may appear in the Storage Server Creation Status wizard screen.

Login credentials or certificate failed message

The following error may appear in the Storage Server Creation Status wizard screen if Cloud Catalyst is being configured for use with a private cloud.



The full error is as follows:

```
Login credentials or certificate verification failed for server
media_server
```

Logs that are related to this error can be found in the
 /usr/opensv/volmgr/debug/tpconfig logs.

Make sure the SSL (Secure Sockets Layer) is configured properly on the private cloud server.

NetBackup supports only Certificate Authority (CA)-signed certificates while it communicates with cloud storage in the SSL mode. Ensure that the cloud server has CA-signed certificate. If it does not have the CA-signed certificate, data transfer between NetBackup and the cloud provider fails in the SSL mode.

The *NetBackup Cloud Administrator's Guide* contains information about SSL mode in private clouds:

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

Failure to add credentials

The wizard may display a failure to add credentials if the media server had previously been configured as a Cloud Catalyst storage server. It also fails if the media server is currently configured as a Cloud Catalyst storage server.

Run the `esfs_cleanup.sh` script on the media server to remove all Cloud Catalyst-related information from the server.

The script is located in the following directory:

`/usr/opensv/esfs/scripts/esfs_cleanup.sh`

Disk pool problems

The following topics present issues concerning disk pools and Cloud Catalyst storage servers.

See [“Disk pool creation problem due to timeout issue”](#) on page 437.

See [“One disk pool for each Cloud Catalyst storage server”](#) on page 438.

Disk pool creation problem due to timeout issue

For disk pool failures while creating a disk pool for a Cloud Catalyst storage server, search for an error message that is similar to the following:

```
Test upload failed, done 5 error code 0, tried 300 times
```

The log is located in `/usr/opensv/netbackup/logs/esfs_filesystem`.

If the error shows that 300 attempts were made, the failure was due to a timeout issue. A slow network connection to the cloud is often the cause. Check network connections before contacting Veritas Support.

One disk pool for each Cloud Catalyst storage server

Each Cloud Catalyst storage server may have only one disk pool.

The following error message appears if the user attempts to create a second disk pool for the cloud storage server:

The Cloud Catalyst disk volume cannot be reconfigured for this storage server.

Problems during cloud storage server configuration

The following errors may appear during configuration of a Cloud Catalyst storage server.

Table 14-10 Cloud storage server configuration errors

Error	Reason
An unknown error occurred in the Cloud Catalyst subsystem	<p>This is a general configuration error.</p> <p>Check the following log for more specific information:</p> <pre>/usr/opensv/esfs/log/ops/esfs-config.log.*</pre> <p>Contact Veritas Support to determine a corrective action.</p>
An error occurred in ESFS initialization	<p>This error may occur if the <code>esfs_init.sh</code> script is run without all of the required parameters.</p> <p>Check the following log for more specific information:</p> <pre>/usr/opensv/esfs/log/ops/esfs-config.log.*</pre>
FUSE not installed	<p>NetBackup installs the Filesystem in Userspace (FUSE) upon installation. If the FUSE version is too old, this message appears.</p> <p>Contact Veritas Support to determine the correct version of FUSE to install it.</p>
Unsupported ESFS operating system	<p>A Cloud Catalyst storage server can be configured on a media server only if the host is running Red Hat Enterprise Linux 7.3 or later.</p> <p>See the NetBackup compatibility lists for support information:</p> <p>http://www.netbackup.com/compatibility</p>
The Cloud Catalyst service could not be started	<p>Check that the <code>vxsfsd</code> daemon has sufficient memory to run.</p> <p>Check the following log for more specific information:</p> <pre>/usr/opensv/esfs/log/ops/esfs-config.log.*</pre>

Table 14-10 Cloud storage server configuration errors (*continued*)

Error	Reason
An ESFS mount could not be found	<p>This error indicates that something is preventing the NetBackup Extendable Storage File System (ESFS) from starting properly. There may also be a problem with the <code>vxesfsd</code> process.</p> <p>Check the following log for more specific information:</p> <pre>/usr/opensv/esfs/log/ops/esfs-config.log. *</pre> <p>Contact Veritas Support to determine a corrective action.</p>

Media server not available in Media Servers tab to add credentials

The Cloud Catalyst storage server can allow access to media servers only if they are Red Hat Enterprise Linux (any version) and have NetBackup 8.1 or later installed. Access is given to media servers by granting them credentials in the **Media Servers** tab of the Cloud Catalyst storage server properties.


If the media server does not have NetBackup 8.1 or later installed, it does not appear in the **Media Servers** tab, so it cannot be selected.

See [“Configuring source control or target control optimized duplication for Cloud Catalyst”](#) on page 277.

Add credentials failed message for media server

The Cloud Catalyst storage server can allow access to media servers only if they are Red Hat Enterprise Linux (any version) and have NetBackup 8.1 or later installed. Access is given to media servers by granting them credentials in the **Media Servers** tab of the Cloud Catalyst storage server properties.

If the media server is non-RHEL, it appears in the **Media Servers** tab. However, if it is selected to have access to the Cloud Catalyst storage server, the operation fails with an **Add credentials failed** message.



See [“Configuring source control or target control optimized duplication for Cloud Catalyst”](#) on page 277.

Status 191: No images were successfully processed

The `no images were successfully processed (191)` message can appear for a duplication job in the **Detailed Status** tab of the **Activity Monitor**.

Storage unit configuration for selecting media servers

Part of creating a cloud storage server includes creating a disk pool and storage unit. When creating the storage unit, you must indicate whether NetBackup can use any available media server or whether NetBackup should use only specific media servers. To configure load balancing for push optimized duplications, select **Only use the following media servers** and then specify the media server that hosts the MSDP storage server.

See [“Configuring source control or target control optimized duplication for Cloud Catalyst”](#) on page 277.

Media write error (84) if due to a full local cache directory

One possible cause for a `media write error (84)` in a Cloud Catalyst environment is described in the following topic.

The administrator configures a local cache directory as part of configuring a Cloud Catalyst storage server. Cache eviction is triggered when data causes this directory to reach the high watermark setting (`HighWatermark`). Once the high watermark is reached, data is purged when the used space reaches the midpoint between `HighWatermark` and `LowWatermark` ($(high+low)/2$) and continues until `LowWatermark` is reached. If the rate of incoming data exceeds the rate where the watermark can be maintained, the jobs begin to fail. In this situation, the directory does not have enough space for new files and the MSDP write operation fails with a `media write error (84)`.

This error is more likely to occur when a relatively small disk is used for demonstration or proof-of-concept testing. To prevent this error in a production environment, make sure that there is at least 4 TB of disk space available for the cache. Make sure that the cache is located on a dedicated file system that is used only for the Cloud Catalyst local cache.

After receiving a `media write error (84)` and addressing the cause, restart all NetBackup services to avoid a `media open error (83)`.

The UNIX `df` command can be used to determine the disk space in the local cache directory. For example:

```
df /esfs
Filesystem      1K-blocks  Used      Available  Use%  Mounted on
/dev/loop0      1007896    1007896         0      100%  /esfs
```

See the following topic for more information about Cloud Catalyst local cache considerations:

See [“About the Cloud Catalyst cache”](#) on page 275.

Troubleshooting restarting ESFS after the Cloud Catalyst storage server is down

If the Cloud Catalyst storage server goes down abnormally, ESFS (`vxesfsd`) may not restart when NetBackup is restarted.

To allow `vxesfsd` to start, check the following location to see if a `rocksd.log` file was created:

```
/cache/fsdb/log
```

Remove the `rocksd.log` file and restart the NetBackup services.

Restarting the vxesfsd process

To restart the Veritas NetBackup Extendable Storage File System Service (`vxesfsd`), the process must first be stopped.

- Before stopping the `vxesfsd` command, verify that `vxesfsd` is not in use by any other process, including the current working directory of any user.
The NetBackup Deduplication Manager (`spad`) and the NetBackup Deduplication Engine (`spoold`) use `vxesfsd`, so they need to be stopped if either is running.
- Before starting the `vxesfsd` command, verify that no data exists in the mount point. If data exists in the mount point, `vxesfsd` fails to restart.

The following methods assume that no jobs are running.

Using the Activity Monitor to stop and start processes specific to vxesfsd

Using the Activity Monitor to stop the processes specific to vxesfsd

- 1 In the **NetBackup Administration Console**, click **Activity Monitor**.
- 2 On the **Daemons** tab, select the `spool`.
- 3 On the **Actions** menu, select **Stop Daemon**.
- 4 Stop `spad` and stop `spool`.
- 5 Stop `vxesfsd`.

Using the Activity Monitor to start the processes specific to vxesfsd

- 1 In the **NetBackup Administration Console**, click **Activity Monitor**.
- 2 On the **Daemons** tab, select `vxesfsd`.
- 3 On the **Actions** menu, select **Start Daemon**.
- 4 Start `spool`.
- 5 Start `spad`.

Using the command line to stop and start processes specific to vxesfsd

To stop and start the processes that are specific to `vxesfsd` (and to avoid stopping all NetBackup processes), run:

```
/usr/opensv/netbackup/bin/nbsvcmon -terminate
/usr/opensv/pdde/pdconfigure/pdde stop
/usr/opensv/esfs/bin/vxesfsd stop
```

To start these processes:

- To start `vxesfsd`, run:

```
/usr/opensv/esfs/bin/vxesfsd start
/usr/opensv/pdde/pdconfigure/pdde start
```

- To start `nbsvcmon`, run:

```
/usr/opensv/netbackup/bin/nbsvcmon &
```

Stopping and starting all NetBackup processes

To stop and start all processes, run the following command on the Cloud Catalyst storage server:

```
/etc/init.d/NetBackup stop
/etc/init.d/NetBackup start
```

Problems restarting vxesfsd

Restarting `vxesfsd` fails if a process accesses the ESFS mount point when NetBackup attempts to shut it down.

See the following log to determine if this issue is the problem:

```
/usr/openv/esfs/log/ops/vxesfs-run.process_id
```

Note that after the hard kill of `vxesfsd`, the target remains busy.

To verify if the ESFS mount point is used by any processes

- 1 Unmount the mount directory:

```
umount mount_dir
```

Or

```
lsof mount_dir
```

- 2 Verify that the directory is not mounted:

```
mount | grep mount_dir
```

Or

```
lsof | grep mount_dir
```

- 3 Restart `vxesfsd`:

```
/usr/openv/esfs/bin/vxesfsd start
/usr/openv/pdde/pdconfigure/pdde start
```

Unable to create CloudCatalyst with a media server that has version earlier to 8.2

If you have upgraded only your NetBackup master server to version 8.2, and want to use CloudCatalyst with a media server that has a version earlier than 8.2, run the following command to enable the cloud web application:

```
/usr/openv/wmc/bin/install/configurePorts -addLegacyCloudService
```

(Or an equivalent path for Windows operating system)

This command adds a new legacy cloud service or connector to Tomcat server's `server.xml` file. Restart the `nbwmc` service.

Cloud Catalyst troubleshooting tools

NetBackup offers several tools that can help to troubleshoot Cloud Catalyst problems.

See [“esfs_cleanup.sh script”](#) on page 444.

See [“esfs_check consistency checking tool”](#) on page 444.

See [“setlsu_iocctl tool”](#) on page 445.

esfs_cleanup.sh script

In certain conditions, use the `esfs_cleanup.sh` script on the Cloud Catalyst storage server to remove all Cloud Catalyst-related information from the host.

For example, when Cloud Catalyst cloud storage is decommissioned. (See [“Decommissioning Cloud Catalyst cloud storage”](#) on page 283.)

The script is located in the following directory:

```
/usr/openv/esfs/scripts/esfs_cleanup.sh
```

Note: Use this script with care as it deletes all data, data loss occurs if used incorrectly. Contact Veritas Technical Support for assistance.

esfs_check consistency checking tool

The `esfs_check` tool determines whether the data that is stored in the cloud and the data in the ESFS database are consistent. This check reports any inconsistencies between the two.

Note: Since the tool reads data from the cloud, it is likely that the NetBackup customer might incur a financial cost by the cloud vendor. Consider contacting Veritas Support before running the tool.

The `esfs_check` tool is located in the following directory: `/usr/openv/esfs/bin`

It uses the following syntax:

```
esfs_check -cachedir path_to_local_cache [-strict] [-verbose]
```

Where:

- `path_to_local_cache` is the path to the local cache directory.

- The `-strict` option provides an additional check to compare the file size in the local cache with the file size in the ESFS database.
- The `-verbose` option enables logging to the following directory:
`/usr/opensv/esfs/log/ops/esfs_check`

Before you run the `esfs_check` tool:

- To ensure accurate results, the system should not have any active backup or duplication jobs to the cloud.
- Shut down NetBackup services except for the ESFS daemons. Only `vxesfsd` and `ocsd` should be running.

Example where the data is consistent

The following example output shows that the data that is stored in the cloud is consistent with the ESFS database.

```
esfs_check -cachedir /root/cache/msdp/ -verbose
consistency check successful
```

Example where the data is inconsistent

The following example output shows that there are inconsistencies between the cloud and the ESFS database.

```
esfs_check -cachedir /root/cache/msdp/ -verbose
File /data/2/97.bin not exist in cloud, expect in cloud
File /data/2/97.bin not exist in cache/cloud
```

Output line 1 means that file `/data/2/97.bin` appears in the ESFS database and should also be in the cloud, but it is not.

Output line 2 means that file `/data/2/97.bin` appears in the ESFS database, but is neither in the cloud nor the local cache.

If the `esfs_check` tool finds an inconsistency, contact Veritas Support.

setlsu_ioctl tool

This tool is used in a disaster recovery scenario to set the cloud bucket name.

The `setlsu_ioctl` tool is located in the following directory: `/usr/opensv/esfs/bin`

It uses the following syntax:

```
./setlsu_ioctl /cache_directory/storage/proc/cloud.lsu name_of_bucket
```

cred_ioctl tool

This tool is used in a disaster recovery scenario to set the cloud credentials.

The `cred_ioctl` tool is located in the following directory: `/usr/opensv/esfs/bin`

The `cred_ioctl` uses this syntax: `./cred_ioctl [-update] <mountPath>
<cloud_username> <cloud_password>`

Unable to obtain the administrator password to use an AWS EC2 instance that has a Windows OS

This error occurs after the instance is launched from an AMI that is converted using automated disaster recovery.

The following error is displayed:

```
Password is not available. This instance was launched from a custom AMI, or the default password has changed. A password cannot be retrieved for this instance. If you have forgotten your password, you can reset it using the Amazon EC2 configuration service. For more information, see Passwords for a Windows Server Instance.
```

For more information, refer to the following articles:

- [Amazon Elastic Compute Cloud Common Messages](#)
- [How to migrate your on-premises domain to AWS Managed Microsoft AD using ADMT](#)

Trouble shooting multi-domain issues

The following sections may help you troubleshoot issues that involve multi-domain scenarios for NetBackup:

See [“Unable to configure OpenStorage server from another domain”](#) on page 446.

See [“MSDP storage server is down when you configure an OpenStorage server”](#) on page 447.

Unable to configure OpenStorage server from another domain

When you try to configure an OpenStorage server from another domain and the error `Login credentials verification failed for server xxxxxx` is displayed, try the following steps to find the root cause:

- Check if the user name and password are correct.

- Check if the NetBackup certificate is deployed to the media server that is used to configure the OpenStorage server. When certificate is not correctly deployed, the following error logs can be found in `pdplugin` log:

```
[ERROR] PDSTS: pd_register: PdvfsRegisterOST(egsusel) failed
(30000:Unknown error 30000)
[ERROR] PDSTS: get_agent_cfg_file_path_for_mount: pd_register()
failed for configuration file:</openv/lib/ost-plugins/egsusel.cfg>
(2060401:UNKNOWN STS ERROR CODE)
```

For more information on using the `nbcertcmd` command to deploy NetBackup certificate for multi-domain, See [“About MSDP multi-domain support”](#) on page 248.

MSDP storage server is down when you configure an OpenStorage server

After configuring an OpenStorage server from another domain if the MSDP storage server is down or unresponsive, run the following steps find the root cause:

- Check if the same MSDP user is used by two or more NetBackup domains.
- Check if there is log entry in `spad.log` as follows:

```
ERR [44] [140589294249728]: 25000: spaProcessing(), It's found that same
msdp user "user1" is used by multiple NBU domains. This is wrong
MultiDomainvconfiguration which will cause potential data loss issue.
Now other NBU domains cannot use msdp user "user1" to access MSDP
services in this server.
```

If there is an error log, the issue is that different NetBackup domains use the same MSDP user to access one MSDP storage server that is not supported by multi-domain.

MSDP server is overloaded when it is used by multiple NetBackup domains

When the MSDP server is used by multiple NetBackup domains and the MSDP server has a high overload, run the following steps to check the workloads from the different domains:

1. Run the following command to get the current tasks status:

For UNIX:

```
/usr/openv/pdde/pdcr/bin/crcontrol --taskstat
```

For Windows:

`<install_path>\Veritas\pdde\crcontrol.exe --taskstat`

2. Check the client column for the list of clients that belong to the NetBackup domain, identify the work load of the clients from one domain, and then work load of one domain.
3. Run the `bpplclients` command on one NetBackup domain to list all clients of that domain.

Migrating to MSDP storage

This appendix includes the following topics:

- [Migrating from PureDisk to the NetBackup MSDP](#)
- [Migrating from another storage type to MSDP](#)

Migrating from PureDisk to the NetBackup MSDP

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

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

Therefore, to replace a NetBackup 5000 series appliance with the NetBackup Media Server Deduplication Option, migrate the storage to a NetBackup **Media Server Deduplication Pool**.

Table A-1 To migrate from PureDisk to NetBackup MSDP

Step	Task	Procedure
Step 1	Install and configure NetBackup	See the <i>NetBackup Installation Guide for UNIX and Windows</i> : http://www.veritas.com/docs/DOC5332

Table A-1 To migrate from PureDisk to NetBackup MSDP (*continued*)

Step	Task	Procedure
Step 2	Configure NetBackup deduplication	See “ Configuring MSDP server-side deduplication ” on page 55.
Step 3	Redirect your backup jobs	Redirect your backup jobs to the NetBackup Media Server Deduplication Pool . See the <i>NetBackup Administrator's Guide, Volume I</i> : http://www.veritas.com/docs/DOC5332
Step 4	Decommission your appliance	Optionally, after the backup images expire on your NetBackup 5000 series appliance, decommission your NetBackup 5000 series appliance. See your NetBackup PureDisk documentation.

See “[Migrating from another storage type to MSDP](#)” on page 450.

Migrating from another storage type to MSDP

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

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

Table A-2 Migrating to NetBackup MSDP

Step	Task	Procedure
Step 1	Configure NetBackup deduplication	See “ Configuring MSDP server-side deduplication ” on page 55.

Table A-2 Migrating to NetBackup MSDP (*continued*)

Step	Task	Procedure
Step 2	Redirect your backup jobs	<p>Redirect your backup jobs to the media server deduplication pool storage unit. To do so, change the backup policy storage destination to the storage unit for the deduplication pool.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>: http://www.veritas.com/docs/DOC5332</p>
Step 3	Repurpose the storage	<p>After all of the backup images that are associated with the storage expire, repurpose that storage.</p> <p>If it is disk storage, you cannot add it to an existing media server deduplication pool. You can use it as storage for another, new deduplication node.</p>

See “[Migrating from PureDisk to the NetBackup MSDP](#)” on page 449.

Index

A

- Active secondary operations 199
- Add at least one index marker 119
- Add credentials failed 277, 439
- adding trusted master server using external certificate 178
- Advanced Encryption Standard (AES) encryption 130
- AES encryption
 - Blowfish encryption 137
- AES-256 134–135
- appliance deduplication 15
- attributes
 - clearing deduplication pool 358
 - clearing deduplication storage server 342
 - OptimizedImage 36
 - setting deduplication pool 352
 - setting deduplication storage server 340
 - viewing deduplication pool 351
 - viewing deduplication storage server 339
- Auto Image Replication
 - about 159
 - and trusted master servers 167
 - Backup operation in source domain 157, 160
 - configuring MSDP replication to a different domain 158
 - for MSDP 184
 - no disk spanning support 160
 - overview 159
 - supported for Cloud Catalyst 258
 - synchronizing clocks of master servers 160
 - Targeted 181
 - targeted 168
 - topology of storage 164

B

- Backup
 - operation 198
- backup
 - client deduplication process 396
- big endian 382, 385
- bpstinfo command 165

byte order 382, 385

C

- cache hits field of the job details 329
- Cannot upload MSDP system files error 437
- capacity and usage reporting for deduplication 330
- case sensitivity
 - in NetBackup names 18
- catalog, MSDP. *See* about recovering the MSDP catalog. *See* MSDP catalog
- certificate, generating for Cloud Catalyst storage server 265
- changing deduplication server hostname 343
- changing the deduplication storage server name and path 343
- CIFS 22
- clearing deduplication pool attributes 358
- client deduplication
 - components 396
 - disabling for a specific client 127
 - host requirements 32
- Client Direct deduplication
 - load balancing 126
 - where not supported 259
- Cloud Catalyst
 - additional drcontrol option 241
 - additional MSDP catalog paths 235
 - Appliance
 - configuring 261
 - documentation 259
 - for deduplication to the cloud 254
 - technical specifications 260
 - direct to cloud (not the preferred use case) 254, 256
 - disaster recovery 289
 - limitations 257–258
 - load balancing 28–29, 85, 126, 258
 - Local cache directory 271
 - local cache directory 268, 275–276
 - logs
 - esfs_filesystem 434

- Cloud Catalyst *(continued)*
 - logs *(continued)*
 - location 434
 - media server version 258, 435
 - overview 254
 - preferred use case 254, 261, 265, 270
 - pull optimized duplication 277
 - push optimized duplication 277
 - vxesfsd process 285
- clustering
 - master server inter-node authentication 181
- Common Internet File System 22
- compacting container files 332
- compression
 - for MSDP backups 128
 - for MSDP optimized duplication and replication 128
 - pd.conf file setting 211
- configuring a deduplication pool 109
- configuring a deduplication storage server 80, 97
- configuring a deduplication storage unit 121
- configuring deduplication 55, 58, 270
- container files
 - about 332
 - compaction 332
 - viewing capacity within 332
- contentrouter.cfg file
 - about 227
 - parameters for data integrity checking 367
 - ServerOptions for encryption 133
- Coordinated Universal Time (UTC) 160
- CR sent field of the job details 329–330
- crcontrol 136
- credentials 34
 - adding NetBackup Deduplication Engine 348
 - changing NetBackup Deduplication Engine 349
 - on media servers in Cloud Catalyst
 - configuration 277–279, 281

D

- data classifications
 - in storage lifecycle policies 192, 198
- data conversion
 - encryption 135
- data integrity checking
 - about deduplication 363
 - configuration settings for deduplication 366
 - configuring behavior for deduplication 364

- data removal process
 - for deduplication 371
- database system error 414
- deactivating media server deduplication 388
- dedup field of the job details 330
- deduplication
 - about credentials 34
 - about fingerprinting 67
 - about the license 51
 - adding credentials 348
 - and Fibre Channel 23
 - and iSCSI 23
 - cache hits field of the job details 329
 - capacity and usage reporting 330
 - changing credentials 349
 - client backup process 396
 - configuration file 209
 - configuring 55, 58, 270
 - configuring optimized synthetic backups 138
 - container files 332
 - CR sent field of the job details 329–330
 - data removal process 371
 - dedup field of the job details 330
 - event codes 431
 - license for 51
 - licensing 52
 - limitations 30
 - media server process 393
 - network interface 35
 - node 19
 - performance 38
 - planning deployment 17
 - requirements for optimized within the same domain 142
 - scaling 41
 - scanned field of the job details 330
 - storage capacity 20
 - storage destination 19
 - storage management 50
 - storage requirements 21
 - storage unit properties 123
 - to the cloud 254
- deduplication configuration file
 - editing 66, 210
 - parameters 211
- deduplication data integrity checking
 - about 363
 - configuring behavior for 364
 - settings 366

- deduplication deduplication pool. *See* deduplication pool
- deduplication disk volume
 - changing the state 359
 - determining the state 359
- deduplication encryption
 - enabling for MSDP backups 132
- deduplication host configuration file 231
 - deleting 232
- deduplication hosts
 - and firewalls 35
 - client requirements 32
 - load balancing server 28
 - server requirements 28
 - storage server 27
- deduplication installation
 - log file 411
- deduplication logs
 - about 407
 - client deduplication proxy plug-in log 409
 - client deduplication proxy server log 409
 - configuration script 410
 - deduplication plug-in log 410
 - NetBackup Deduplication Engine 411
 - NetBackup Deduplication Manager 412
 - VxUL deduplication logs 407
- deduplication node
 - about 19
 - adding a load balancing server 206
 - removing a load balancing server 345
- deduplication optimized synthetic backups
 - about 36
- deduplication plug-in
 - about 391
 - log file 410
- deduplication plug-in configuration file
 - configuring 61
- deduplication pool. *See* deduplication pool
 - about 107
 - changing properties 353
 - clearing attributes 358
 - configuring 109
 - deleting 361
 - determining the state 350
 - properties 116
 - setting attributes 352
 - viewing 350
 - viewing attributes 351
- deduplication port usage
 - about 35
 - troubleshooting 426
- deduplication processes do not start 417
- deduplication rate
 - how file size affects 39
- deduplication reference database
 - about 391
- deduplication registry
 - resetting 232
- deduplication servers
 - components 390
 - host requirements 28
- deduplication storage capacity
 - about 20
 - viewing capacity in container files 332
- deduplication storage destination 19
- deduplication storage requirements 21
- deduplication storage server
 - about 27
 - change the name 343
 - changing properties 341
 - clearing attributes 342
 - components 390
 - configuration failure 416
 - configuring 80, 97
 - defining target for Auto Image Replication 164
 - deleting 346
 - deleting the configuration 347
 - determining the state 338
 - editing configuration file 229
 - getting the configuration 228
 - recovery 381
 - replacing the host 385
 - setting attributes 340
 - setting the configuration 230
 - viewing 338
 - viewing attributes 339
- deduplication storage server configuration file
 - about 227
- deduplication storage server name
 - changing 343
- deduplication storage type 19
- Deduplication storage unit
 - Only use the following media servers 123
 - Use any available media server 123
- deduplication to the cloud
 - direct to cloud 256

- deduplication, compression rate
 - monitoring 325
- deleting backup images 361
- deleting Cloud Catalyst cloud storage 283
- deleting deduplication host configuration file 232
- df command 428, 441
- disaster recovery
 - for Cloud Catalyst 289
 - protecting the data 45
 - recovering the storage server after catalog recovery 384
- disk failure
 - deduplication storage server 380
- disk logs 334
- disk logs report 332
- disk pool status report 331, 334
- disk pools
 - cannot delete 422
 - changing the state 351
 - creation failure for Cloud Catalyst 437
 - failure for Cloud Catalyst 438
- disk storage unit report 334
- Disk type 123
- disk volume
 - changing the state 359
 - determining the state of a deduplication 359
 - volume state changes to down 420
- domains
 - replicating backups to another. *See* Auto Image Replication
- DownloadThreads Cloud Catalyst setting 274
- Duplication
 - optimized 199
- duplication jobs, cancelling 423
- Duplication to remote master. *See* Auto Image Replication

E

- Enable file recovery from VM backup 200
- encryption
 - enabling for MSDP backups 132
 - pd.conf file setting 214
 - SHA-2 67, 134–135
- endian
 - big 382, 385
 - little 382, 385
- ESFS database 434, 444
- esfs.json configuration file 271, 274
- esfs_check tool 444

- esfs_cleanup.sh script 284, 437, 444
- event codes
 - deduplication 431

F

- Fibre Channel
 - and iSCSI comparison 23
 - support for 23
- Fibre Channel and iSCSI comparison for MSDP 23
- file system
 - CIFS 22
 - NFS 22
 - Veritas File System for deduplication storage 50
 - ZFS 22
- File System Database (FSDB) 284–285
- Filesystem in Userspace (FUSE) 284–285, 438
- fingerprinting
 - about deduplication 67
- firewalls and deduplication hosts 35
- FlashBackup policy
 - Maximum fragment size (storage unit setting) 123
- FQDN or IP Address property in Resilient Network host properties 202
- FSDB (File System Database) 284–285
- FUSE (Filesystem in Userspace) 284–285, 438

G

- garbage collection. *See* queue processing
- Global logging level property 272

H

- HighWatermark Cloud Catalyst setting 273, 435, 440
- host requirements 28

I

- I AM Role 270
- images on disk report 334
- Import
 - operation 193
- iSCSI
 - and Fibre Channel comparison 23
 - support for 23

J

- job ID search in unified logs 405

K

KMS encryption 262, 266, 268

L

legacy logging 406

Level (logging) Cloud Catalyst setting 272

license

for deduplication 51

license information failure

for deduplication 416

licensing deduplication 52

limitations

media server deduplication 30

little endian 382, 385

load balancing server

about 28

adding to a deduplication node 206

for deduplication 28

removing from deduplication node 345

Local cache directory for Cloud Catalyst 266, 268, 271, 428

logging

legacy 406

level default for Cloud Catalyst 272

logs

about deduplication 407

Auto Image Replication 412

client deduplication proxy plug-in log 409

client deduplication proxy server log 409

deduplication configuration script log 410

deduplication installation 411

deduplication plug-in log 410

disk 334

NetBackup Deduplication Engine log 411

NetBackup Deduplication Manager log 412

optimized duplication 412

VxUL deduplication logs 407

LowWatermark Cloud Catalyst setting 273

M

maintenance processing. *See* queue processing
master servers

inter-node authentication for clustering 181

Maximum concurrent jobs 124

Maximum fragment size 123

Media open error (83) 441

media server deduplication

process 393

Media Server Deduplication Pool 107, 184

changing the state 351

creating directories for 250-TB support 118

enable 250-TB support 76

media server deduplication pool. *See* deduplication pool

Media write error (84) 440

migrating from PureDisk to NetBackup

deduplication 449

migrating to NetBackup deduplication 450

mklogdir.bat 406

MSDP

replication 156

replication target, configuring 184

MSDP catalog 233, 377

See also MSDP catalog backup

See also MSDP catalog recovery

about the catalog backup policy 234

about the shadow catalog 233

changing the number of shadow copies 238

changing the shadow catalog path 236

changing the shadow catalog schedule 237

shadow copy log files 409

MSDP catalog backup

about protecting the MSDP catalog 234

configuring 239, 244

MSDP catalog recovery

about 377

error codes 429

process the transaction queue. 362

recover from a shadow copy 378

MSDP drcontrol utility

options 240

MSDP replication

about 37

MSDP storage rebasing. *See* rebasing

mtstrm.conf file

configuring 61

N

nbstserv process 194

NDMP

storage units 199

NetBackup

naming conventions 18

NetBackup 5200 series appliance

as a storage destination 20

NetBackup 5300 series appliance

as a storage destination 20

- NetBackup appliance deduplication 15
- NetBackup Appliance Shell Menu 263
- NetBackup Cloud Catalyst 254
 - See also* Cloud Catalyst
- NetBackup deduplication
 - about 14
 - license for 51
- NetBackup Deduplication Engine
 - about 391
 - about credentials 34
 - adding credentials 348
 - changing credentials 349
 - logs 411
- NetBackup Deduplication Manager
 - about 391
 - logs 412
- NetBackup deduplication options 14
- NetBackup Service Monitor (nbsvcmon) 442
- network interface
 - for deduplication 35
- NFS 22
- node
 - deduplication 19

O

- OpenStorage
 - optimized duplication 199
- OpenStorage Disk Option 192
- OpsCenter 258
- optimized deduplication
 - configuring bandwidth 191
 - configuring for MSDP 149
 - limitations 142
 - logs 412
 - pull configuration within the same domain 147
 - separate network for 140
- optimized deduplication copy
 - guidance for 144
- optimized duplication
 - about 37
 - about the media server in common within the same domain 144
 - load balancing for 264, 266
 - push and pull considerations for Cloud Catalyst 277
 - push configuration within the same domain 144
- optimized duplication encryption
 - configuring for MSDP 133

- optimized MSDP deduplication
 - requirements 142
 - within the same domain 141
- optimized synthetic backups
 - configuring for deduplication 138
 - deduplication 36
- OptimizedImage attribute 36

P

- pd.conf file
 - about 209
 - editing 66, 210
 - parameters 211
- pdde-config.log 410
- performance
 - deduplication 38
 - monitoring deduplication rate 325
- port usage
 - and deduplication 35
 - troubleshooting 426
- Postponed secondary operations 199
- Priority for secondary operations 198
- provisioning the deduplication storage 47
- pull optimized duplication for Cloud Catalyst 277
- PureDisk Deduplication Pool 107
- push optimized duplication for Cloud Catalyst 277

Q

- queue processing 362
 - invoke manually 362

R

- ReadOnly Cloud Catalyst setting 274
- rebasing
 - about 368
 - FP_CACHE_PERIOD_REBASING_THRESHOLD parameter 217
 - FP_CACHE_REBASING_THRESHOLD parameter 217
 - RebaseMaxPercentage parameter 371
 - RebaseMaxTime parameter 371
 - RebaseMinContainers parameter 371
 - RebaseScatterThreshold parameter 371
 - server-side rebasing parameters 370
- recovery
 - deduplication storage server 381
 - from deduplication storage server disk failure 380

- Red Hat Linux
 - deduplication processes do not start 417
 - replacing the deduplication storage server 385
- replication
 - between NetBackup domains. *See* Auto Image Replication
 - configuring MSDP replication to a different domain 158
 - for MSDP 37, 156
 - to an alternate NetBackup domain. *See* Auto Image Replication
- replication encryption
 - configuring for MSDP 133
- reports
 - disk logs 332, 334
 - disk pool status 331, 334
 - disk storage unit 334
- resetting the deduplication registry 232
- Resiliency property in Resilient Network host properties 202
- Resilient connection
 - Resilient Network host properties 200
- resilient network connection
 - log file 413
- Resilient Network host properties 200
 - FQDN or IP Address property in 202
 - Resiliency property in 202
- restores
 - at a remote site 374
 - how deduplication restores work 373
 - specifying the restore server 375
- reverse host name lookup
 - prohibiting 416
- reverse name lookup 416
- rolling conversion
 - AES encryption 134

S

- scaling deduplication 41
- scanned field of the job details 330
- Secure Hash Algorithm 67, 134
- server not found error 415
- setlsu_ioctl tool 445
- setting deduplication pool attributes 352
- SHA-2 67, 134–135
- SHA-512/256 134–135
- SLP Parameters 155
- snapshots
 - operation type 198
- spa.cfg file
 - parameters for data integrity checking 368
- SSE4.2 259
- status code 191 (no images were successfully processed) 440
- storage capacity
 - about 20
 - for deduplication 20
 - viewing capacity in container files 332
- Storage Lifecycle Manager service (nbstserv) 194
- storage lifecycle policies
 - Active secondary operations 199
 - best practices document 192
 - cancelling duplication jobs 423
 - Data classification setting 198
 - hierarchy 196
 - operations 194
 - Postponed secondary operations 199
 - Priority for secondary operations 198
 - Storage lifecycle policy name 198
 - utility 192
 - Validate Across Backup Policies button 199
- storage paths
 - about reconfiguring 343
 - changing 343
- storage rebasing. *See* rebasing
- storage requirements
 - for deduplication 21
- storage server
 - about the configuration file 227
 - change the name 343
 - changing properties for deduplication 341
 - changing the name 343
 - components for deduplication 390
 - configuring for deduplication 80, 97
 - deduplication 27
 - define target for Auto Image Replication 164
 - deleting a deduplication 346
 - deleting the deduplication configuration 347
 - determining the state of a deduplication 338
 - editing deduplication configuration file 229
 - getting deduplication configuration 228
 - recovery 381
 - replacing the deduplication host 385
 - setting the deduplication configuration 230
 - viewing 338
 - viewing attributes 339
- storage server configuration
 - getting 228

- storage server configuration *(continued)*
 - setting 230
- storage server configuration file
 - editing 229
- storage type
 - for deduplication 19
- storage unit
 - configuring for deduplication 121
 - properties for deduplication 123
 - recommendations for deduplication 124
- storage unit groups
 - and storage lifecycle policies 199
 - not supported for Auto Image Replication
 - source 157, 160
- Storage unit name 123
- Storage unit type 123
- stream handlers
 - NetBackup 39
- synthetic backups
 - no Auto Image Replication support 160

T

- Targeted A.I.R. 181
- throttling, configuring for Cloud Catalyst 435
- topology of storage 164–165
- troubleshooting
 - database system error 414
 - deduplication backup jobs fail 417
 - deduplication processes do not start 417
 - general operational problems 421
 - host name lookup 416
 - installation fails on Linux 413
 - no volume appears in disk pool wizard 416
 - server not found error 415
- trusted master servers
 - adding 173
 - for Auto Image Replication 167
 - removing 179

U

- unified logging 401
 - format of files 402
- uninstalling media server deduplication 388
- UploadThreads Cloud Catalyst setting 274
- UTC (Coordinated Universal Time) 160

V

- Validate Across Backup Policies button in SLP 199

- Validation Report tab 199
- Veritas NetBackup Extendable Storage File System
 - Service (vxesfsd) 271, 285, 441, 443, 445
- Veritas NetBackup Extendable Storage Proxy
 - (vxesp) 445
- viewing deduplication pool attributes 351
- viewing storage server attributes 339
- VM backup 200
- volume manager
 - Veritas Volume Manager for deduplication
 - storage 50
- VxESFS (Veritas NetBackup Extendable Storage File System) 271
- vxesfsd (Veritas NetBackup Extendable Storage File System Service) 285, 441, 443, 445
- vxlogview command 402
 - with job ID option 405

Z

- ZFS 22