

NetBackup™ Snapshot Client Administrator's Guide

UNIX, Windows, and Linux

Release 10.0



NetBackup™ Snapshot Client Administrator's Guide

Document version: 10.0

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

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

Legal Notice

Copyright © 2022 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>

Contents

Chapter 1	Introduction	12
	About Snapshot Client	12
	About snapshot management	13
	About snapshots	13
	About off-host backup support	13
	About Instant Recovery	15
	About snapshot methods for disk arrays	15
	About block level incremental backup	15
	About snapshot basics	15
	About the copy-on-write snapshot type	16
	About the mirror snapshot type	16
	Benefits of copy-on-write versus mirror	17
	About local backup of a snapshot	18
	Off-host backup overview	19
	About file and volume mapping methods	20
	Off-host backup methods	20
	About alternate client backup	21
	NetBackup media server data mover example (UNIX only)	26
	Snapshot Client requirements	28
	Snapshot Client restrictions	28
	Snapshot Client terminology	29
	Snapshot Client assistance	32
	About open file backups for Windows	33
 Chapter 2	 Installation	 34
	Installation prerequisites for Snapshot Client	34
	Snapshot Client installation notes	34
	About distributing Client Software in mixed-platform environments	36
	About creating log directories	36
	About the snapshot state file	36

Chapter 3	Policy configuration	38
	Notes on Snapshot Client policies	38
	Configuring a Snapshot Client policy	39
	Backup Selections tab options when configuring a policy	42
	Off-host backup configuration options	42
	Selecting the snapshot method	43
	Configuration parameters for Snapshot Client	45
	Snapshot Resources	46
	Configuring backup scripts	46
	About using alternate client backup	47
	Alternate client backup requirements	47
	Configuring alternate client backup	48
	Example alternate client backup configurations	49
	Policy configuration tips	49
	Maximum pathname length	49
	Snapshot tips	49
	Multiple data streams	50
	About disabling snapshots	50
	Disabling Snapshot Client snapshots	50
Chapter 4	NAS snapshot configuration	52
	About NAS snapshot overview	52
	Dynamic data streaming for D-NAS workloads	53
	Understanding the features	54
	Dynamic streaming parameters	55
	Prerequisites for NetBackup DNAS configuration	56
	Configuring dynamic data streaming with backup host pool	57
	Configuring a backup host pool	58
	Discover plugins	59
	Minimum supported backup host versions for different features	62
	Limitations and considerations	62
	Setting up a NAS-Data-Protection policy	64
	About restores from NAS-Data-Protection policy	69
Chapter 5	FlashBackup configuration	71
	About FlashBackup	71
	FlashBackup restrictions	72
	Restores of Windows encrypted files and hard links	73
	Configuring a FlashBackup policy	73
	Configuring FlashBackup policy for backward compatibility (UNIX only)	77

	About the cache partition	79
	Requirements for the cache partition	80
	Directives for multiple data streams	80
Chapter 6	Instant Recovery configuration	83
	About Instant Recovery capabilities	83
	Instant Recovery requirements	84
	Instant Recovery restrictions	84
	Giving full server privileges to the media server	84
	About Instant Recovery	85
	About snapshot and backup for Instant Recovery	86
	About NetBackup catalog maintenance	86
	About snapshot management	87
	Means of controlling snapshots	87
	Configuring a policy for Instant Recovery	87
	About sizing the cache for Instant Recovery copy-on-write snapshots	88
	Cache size during restore	89
	Setting an adequate size for the snapshot cache	89
	Instant Recovery for databases	89
	About storage lifecycle policies for snapshots	90
	Configuring a storage lifecycle policy to manage snapshot-based backups for Instant Recovery	90
	Storage lifecycle policies and Snapshot Client troubleshooting	92
Chapter 7	Configuration of software-based snapshot methods	93
	Software-based snapshot methods	93
	About nbu_snap	93
	Cache device requirements	94
	About VxFS_Checkpoint	98
	About VxFS_Snapshot	100
	About VxVM	100
	About FlashSnap	103
	About VVR	106
	About NAS_Snapshot	108
	About VSS	108

Chapter 8	Support for Cluster Volume Manager Environments (CVM)	111
	About support for CVM environments	111
	Note on NetBackup and CVM	112
	About enabling VxVM or FlashSnap snapshots in a CVM environment	112
	About enabling the NetBackup client to execute VxVM commands on the CVM master node	113
Chapter 9	Configuration of snapshot methods for disk arrays	114
	OS-specific configuration tasks	114
	About dynamic multi-pathing	114
	HBA configuration	115
	Verifying NetBackup client access, zoning, and LUN masking	115
	About IBM DS6000 and DS8000 arrays	116
	IBM DS6000 and DS8000 software requirements	116
	Preconfiguration for IBM arrays	116
	Configuring NetBackup to access the IBM DS6000 or DS8000 array	117
	Configuring the IBM array for NetBackup	117
	Using DSCLI commands to obtain unique IBM identifiers	118
	Configuring a NetBackup policy for IBM_DiskStorage_FlashCopy	121
	For further reference on IBM arrays	122
	About IBM DS4000 array	122
	Array preconfiguration tasks	122
	IBM 4000 software requirements	123
	Verifying NetBackup client access, zoning, and LUN masking	123
	Configuring NetBackup to access the IBM DS4000 array	125
	Configuring the IBM 4000 array for NetBackup	125
	Configuring a NetBackup policy for IBM_StorageManager_FlashCopy	126
	About Hitachi SMS/WMS/AMS, USP/NSC, USP-V/VM	127
	Hitachi array software requirements	127
	Preconfiguration for Hitachi	128
	About communication between NetBackup and the Hitachi array	128
	Determining if the Hitachi command devices are visible	128

About configuring the Hitachi array for NetBackup	129
Obtaining the Hitachi array serial number and the unique device identifiers	129
Configuring a NetBackup policy for Hitachi_ShadowImage or Hitachi_CopyOnWrite	130
About HP-XP arrays	132
HP-XP array software requirements	132
Preconfiguration for HP-XP	132
About communication between NetBackup and the HP-XP array	133
Determining if the HP-XP command devices are visible	133
About configuring the HP-XP array for NetBackup	133
Obtaining the array serial number and unique HP-XP identifiers	133
Configuring a NetBackup policy for HP_XP_BusinessCopy and HP_XP_Snapshot	134
About array troubleshooting	135
Troubleshooting issues pertaining to all arrays	135
Troubleshooting NetBackup and EMC CLARiiON arrays	136
Troubleshooting NetBackup and EMC Symmetrix arrays	138
Troubleshooting NetBackup and HP EVA arrays	139
Troubleshooting IBM DS6000 and DS8000 arrays	140
Troubleshooting IBM4000 arrays	143
Troubleshooting Hitachi arrays	145

Chapter 10 Notes on Media Server and Third-Party Copy methods 152

Disk requirements for Media Server and Third-Party Copy methods	152
Directives for Media Server and Third-Party Copy methods	153
Storage units for Media Server and Third-Party Copy methods	153
Preventing multiplexing on a third-party copy backup	153
Raw partition backups	153
Increasing the client read timeout for all clients	154
Further information on off-host data mover backups	154

Chapter 11 Backup and restore procedures 155

About performing a backup	155
About performing a restore	156
Instant Recovery restore features	156
About Instant Recovery: point in time rollback	156

Notes for restoring individual files from an Instant Recovery snapshot	162
About configurations for restore	163
About restoring over the LAN	163
About restoring over the SAN to a host acting as both client server and media server	163
About restoring directly from a snapshot	164
About restoring from a disk snapshot	165
About restoring on UNIX	165
About restoring on Windows	168
 Chapter 12 Snapshot management	171
About NetBackup and CloudPoint integration for snapshot management	171
Troubleshooting Snapshot Client and NetBackup integration issues	173
 Chapter 13 Troubleshooting	176
About gathering information and checking logs	177
Logging directories for UNIX platforms	178
UNIX logging directories for backup	178
UNIX logging directories for restore	179
snapctl driver messages	179
Logging folders for Windows platforms	179
Windows logging folders for backup	180
Windows logging folders for restore	180
Configuring VxMS logging	181
Customer support contact information	183
Latest patches and updates	184
Snapshot provider information	184
Important notes on &newProductName;	184
&newProductName; installation problems	186
FlashBackup and status code 13	186
The FlashBackup cache partition may have run out of space	186
Removing stale snapshots (Solaris)	187
Single file restore from a FlashBackup Instant Recovery snapshot of a file protected by Windows VSS writer fails	187
Identifying and removing a left-over snapshot	188
Removing a VxVM volume clone	193
Alternate client restore and backup from a snapshot fails	196
Restore from a snapshot fails with status 2800	196

Raw Partition restore fails with the message 'FlashBackup-Windows policy restore error'	197
Snapshot creation fails with error 156	197
Snapshot fails with error 20	197
Restore from a snapshot fails with status 133	198
Backup from snapshot fails with error 50	198
Backup from snapshot parent job fails with error 4213: Snapshot import failed	198
Snapshot job fails and the snapshot command does not recognize the volume name	199
Snapshot creation fails with error 4220	199
Snapshot creation fails when the same volume is mounted on multiple mount points of the same host	200
Snapshot-based backup and restore failure	201
Multiple snapshot jobs fail with code 156 or 1541.	201
FlashBackup policy fails, with multiple backup selections [Cache =]	201
Partial backup failure with 'Snapshot encountered error 156'	202
Backup of file system validation fails with error 223	202
Policy validation fails if the specified CIFS share path contains a forward slash	202
Policy validation fails with status code 223	203
An NDMP snapshot policy for wildcard backup fails with error 4201	203
Troubleshooting with bpfis log	203
Limitations of using HP-UX 11.31	203
While configuring a policy for snapshot operations, the NetBackup Administration Console hangs	204
Accelerator enabled incremental backup of NetApp NAS volume	204
Snapshot method: Auto	205
Appendix A Managing nbu_snap (Solaris)	206
About managing nbu_snap	206
Cache for nbu_snap	206
About determining cache size	207
About terminating nbu_snap	207
nbu_snap commands	207
Appendix B Overview of snapshot operations	212
Introduction to snapshot operations	212
Pre and post snapshot creation operations	213
About quiescing the system	213

About quiescing the database application	214
About quiescing the stack	215
File system quiesce	215
Volume manager data caching	215
How copy-on-write works	215
Index	220

Introduction

This chapter includes the following topics:

- [About Snapshot Client](#)
- [About snapshot management](#)
- [About snapshot basics](#)
- [Benefits of copy-on-write versus mirror](#)
- [About local backup of a snapshot](#)
- [Off-host backup overview](#)
- [Off-host backup methods](#)
- [Snapshot Client requirements](#)
- [Snapshot Client terminology](#)
- [Snapshot Client assistance](#)
- [About open file backups for Windows](#)

About Snapshot Client

Table 1-1 Snapshot Client features at a glance

Snapshot Client feature	Description
Snapshot	
Instant Recovery	Makes the backups available for recovery from snapshot.

Table 1-1 Snapshot Client features at a glance (*continued*)

Snapshot Client feature	Description
Off-host backup	Shifts the burden of backup processing onto a separate backup agent, reducing the backup effect on the client's computing resources. The backup agent sends the client's data to the storage device. Includes alternate the client, data mover, and virtual machine host (for VMware).
NetBackup for VMware	Backs up and restores Windows and Linux VMware virtual machines (guest operating systems). See the NetBackup for VMware Administrator's Guide .
NAS snapshot	Makes snapshot-based backup of data on a Network Attached Storage (NAS) host.
Block level incremental backup (BLIB)	Enables NetBackup to back up only the changed data blocks of VMware virtual machines and Oracle or DB2 database files.

About snapshot management

These topics describe the features of Snapshot Client.

About snapshots

A snapshot is a point-in-time, read-only, disk-based copy of a client volume. After the snapshot is created, NetBackup backs up data from the snapshot, not directly from the client's primary or original volume. Users and client operations can access the primary data without interruption while data on the snapshot volume is backed up. The contents of the snapshot volume are cataloged as if the backup was produced directly from the primary volume. After the backup is complete, the snapshot-based backup image on storage media is indistinguishable from a traditional, non-snapshot backup image.

All the features of Snapshot Client (including off-host backup, FlashBackup, and Instant Recovery) require the creation of a snapshot.


About off-host backup support

Another major component of NetBackup Snapshot Client is support for off-host backup. Off-host backup shifts the burden of backup processing onto a separate

backup agent, greatly reducing the backup effect on the client's computing resources. The backup agent sends the client's data to the storage device.

Figure 1-1 shows a backup agent.

Figure 1-1 Backup agent for off-host backup



The backup agent can be any of the following:

- An additional (alternate) client
- A NetBackup media server or a third-party copy device that implements the SCSI Extended Copy command
- A NAS host (Network Attached Storage)

Note that many types of devices are designed to act as third-party copy devices, such as routers, bridges, robotic libraries, and disk arrays. The backup agent can direct the data to SCSI-attached storage or to storage on the SAN.

Note: NetBackup supports off-host backup of Oracle database in the SFRAC (Storage Foundation Real Application Clusters) environment. For more details, refer to the [NetBackup for Oracle System Administrator's Guide](#).

About Instant Recovery

This feature makes backups available for quick recovery from disk. Instant Recovery combines snapshot technology—the image is created with minimal interruption of user access to data—with the ability to do rapid snapshot-based restores. The snapshot is retained on disk as a full backup image. The snapshot can also be the source for an additional backup copy to tape or other storage.

Instant Recovery makes possible three additional variations of restore: block-level restore, file promotion, and snapshot rollback.

See [“Instant Recovery restore features”](#) on page 156.

About snapshot methods for disk arrays

Snapshot Client supports snapshot methods for many disk arrays. Each array method is designed for a particular disk array series and a particular snapshot type.

These methods support the following:

- Auto, Mirror, clone, and copy-on-write snapshot types.
- NetBackup Instant Recovery, both snapshot-only and snapshot-to-tape backup.
- Local backup or alternate client backup.
- Backup of Oracle, Exchange, and DB2 database clients.

About block level incremental backup

Block level incremental backup enables NetBackup to back up only the changed data blocks of VMware virtual machines and Oracle or DB2 database files. For details, refer to the appropriate NetBackup database agent guide or to the [NetBackup for VMware Administrator's Guide](#).

About snapshot basics

The large active databases or file systems that must be available around-the-clock are difficult to back up without incurring a penalty. Often, the penalty takes one of the following forms:

- To allow time for the backup, the entire database is taken offline or the file system is unmounted. The result is suspension of service and inconvenience to users.
- The copy is made very quickly but produces an incomplete version of the data, since some transactions have failed to complete.

A solution to this problem is to create a snapshot of the data. A snapshot captures the data at a particular instant, without causing the client downtime. The resulting capture or snapshot can be backed up without affecting the performance or availability of the file system or database. Without a complete, up-to-date snapshot of the data, a correct backup cannot be made.

When a NetBackup media server or third-party copy device manages the backup, the data to back up must be contained in a snapshot. The backup agent can only access the data by means of the raw physical disk. Once the data is captured as a snapshot, the NetBackup client "maps" the logical representation of the data to its physical disk address. These disk addresses are sent to the backup agent over the LAN. The data is then read from the appropriate disk by the backup agent.

See ["Off-host backup overview"](#) on page 19.

Two types of snapshots are available, both supported by NetBackup: copy-on-write and mirror (or clone). In **Auto** mode, NetBackup automatically decides the best type of snapshot which the storage array supports.

About the copy-on-write snapshot type

A copy-on-write type of snapshot is a detailed account of data as it existed at a certain moment. A copy-on-write snapshot is not a copy of the data, but a specialized account of it.

A copy-on-write snapshot is created in the client's file system or in a raw partition. The copy-on-write is not created as a complete copy of the client data on a separate or a mirror disk. The snapshot is then backed up to storage as specified in the backup policy. Users can access their data without interruption, as though no backup is underway. The file system is paused long enough to assemble a transnationally consistent record.

See ["How copy-on-write works"](#) on page 215.

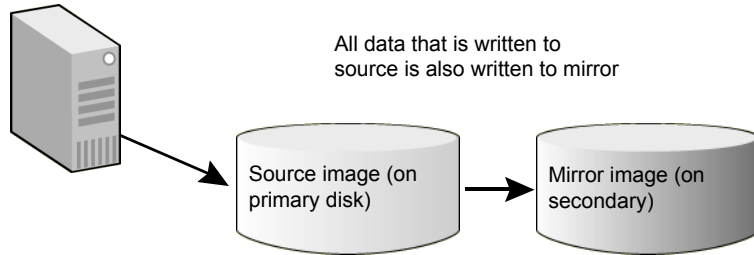
See ["Benefits of copy-on-write versus mirror"](#) on page 17.

About the mirror snapshot type

Unlike a copy-on-write, a mirror is a complete data copy stored on a separate disk, physically independent of the original. Every change or write to the data on the primary disk is also made to the copy on the secondary disk. The copy is a "mirror" image of the source data.

Figure 1-2 Source is copied to mirror

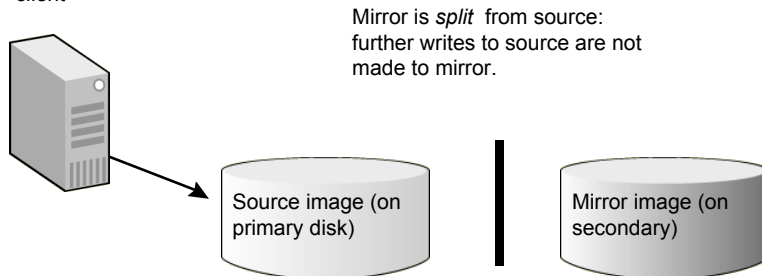
NetBackup
client



As in a copy-on-write, transactions are allowed to finish and new I/O on the primary disk is briefly halted. When the mirror image is brought up-to-date with the source, the mirror is split from the primary. After the mirror is split, new changes can be made to the primary but not to the mirror. The mirror can now be backed up (see next diagram).

Figure 1-3 Mirror is split from source

NetBackup
client



If the mirror is to be used again it must be brought up-to-date with the primary volume (synchronized). During synchronization, the changes that were made to the primary volume—while the mirror was split—are written to the mirror.

Since mirroring requires a complete copy of the primary on a separate device (same size as the primary), it consumes more disk space than copy-on-write.

See [“Benefits of copy-on-write versus mirror”](#) on page 17.

Benefits of copy-on-write versus mirror

[Table 1-2](#) compares the benefits of the two types of snapshots.


Table 1-2 Comparative benefits of copy-on-write and mirror

Benefits of copy-on-write	Benefits of mirror
■ It consumes less disk space: No need for secondary disks containing complete copies of source data. ■ Relatively easy to configure (no need to set up mirror disks). ■ Creates a snapshot much faster than one created by a large, unsynchronized mirror, because mirror synchronization can be time consuming.	■ It has less effect on the performance of the host being backed up (NetBackup client), because the copy-on-write mechanism is not needed. ■ Allows for faster backups: The backup process reads data from a separate disk (mirror) operating independently of the primary disk that holds the client's source data. Unlike copy-on-write, disk I/O is not shared with other processes or applications. Apart from NetBackup, no other applications have access to the mirror disk. During a copy-on-write, other applications as well as the copy-on-write mechanism can access the source data. Note: If additional disk drives are available and volumes have been configured with the Veritas Volume Manager, a mirror snapshot method is usually a good choice.

About local backup of a snapshot

A snapshot can be backed up to any NetBackup storage device. A Fibre Channel network or SAN is not required. The following diagram shows a network configuration sufficient for backing up a snapshot on the primary client (sometimes referred to as a "local" snapshot backup). The network configuration is identical to the configuration for normal NetBackup (no snapshot).

Figure 1-4 Snapshot backup on local network (no Fibre Channel/SAN required)



The figure shows the following phases in the local backup process:

Phase	Action
Phase 1	NetBackup master server tells the client to create the snapshot of the primary data.
Phase 2	Client reads the snapshot, formats a backup image reflecting the primary data, and writes the image to the media server.
Phase 3	Media server reads the backup image.
Phase 4	Media server writes data to local storage.

Off-host backup overview

One of the principal goals of NetBackup Snapshot Client is to move I/O processing off the primary NetBackup client to a backup agent.

[Table 1-3](#) describes the types of backup agents.

Table 1-3 Type of backup agents

Type of backup agent	Description
Alternate client	A secondary or an alternate client performs the backup on behalf of the primary client. Compared to the other off-host methods, this approach reduces the backup I/O burden on the primary client to the greatest extent.
Data mover: NetBackup media server (UNIX clients only)	A NetBackup media server reads raw data from the client snapshot and writes it to a storage device, using mapping information that the client provides.
Data mover: Network Attached Storage	An NDMP (NAS) host performs the snapshot-only backup for Instant Recovery only.
Data mover: NDMP	Use to replicate NDMP snapshots. Select this agent in a policy that uses NDMP with Replication Director. For more information about configuring a policy to use NDMP with Replication Director, see the NetBackup Replication Director Solutions Guide .
Virtual machine host	A VMware backup host performs backups on behalf of the virtual machines. The host can also be configured as a NetBackup master or media server.

About file and volume mapping methods

The NetBackup media server and Third-Party Copy Device backup agents are unaware of logical organizations of data such as file systems and volume managers. The agent can access the data only from the physical disk address location. In order for NetBackup to perform this type of backup, it must translate the logical representation of the data to its physical disk addresses. This logical-to-physical translation process is referred to as mapping the data. During the backup, the mapping information is transmitted to the media server.

The mapping methods are installed as part of the NetBackup Snapshot Client product. Depending on whether the backup data is configured over physical devices, logical volumes, or file systems, NetBackup automatically selects the correct mapping method.

Off-host backup methods


NetBackup Snapshot Client supports several forms of off-host backup, as explained in the following sections.

About alternate client backup

With this feature, all backup processing is off-loaded to another client. Off-loading the work to an alternate client saves computing resources on the primary client. The alternate client handles the backup I/O processing, and the backup has little or no effect on the primary client.

The following diagram shows alternate client backup. A NetBackup master server is connected by means of a local or wide-area network to two clients and a media server. The primary NetBackup client contains the data to be backed up. A snapshot of that data is created on the alternate client (perhaps at another site). The alternate client creates a backup image from the snapshot, using original path names, and streams the image to the media server.

Figure 1-5 Alternate client backup: backup is performed on alternate client



The figure shows the following phases in the alternate client backup process:

Phase	Action
Phase 1	Primary and alternate client collaborates to create the snapshot on the alternate client.
Phase 2	Alternate client sends the snapshot data to the media server.
Phase 3	Media server reads the snapshot data from the alternate client.
Phase 4	Media server writes data to local storage.

About data sharing between clients


For alternate client backup, the original (primary) and alternate client must collaborate to create a snapshot. The following sections highlight two configurations: hardware array split-mirror snapshot and volume manager snapshot with data replication. Other configurations are possible, depending on the hardware and the snapshot method.

Alternate client backup split mirror examples

The alternate client has access to the mirror disks that contains a snapshot of the primary client's data. Before the backup, the mirror is split from the primary disk, which creates the snapshot on the mirror disk. The alternate client has access to the mirror disk, from which it creates and streams a snapshot-based backup image to the media server. After the backup, the mirror can be optionally resynchronized with the primary disk.

Note: The mirror disk need not be visible to the primary client, only to the alternate client.

Figure 1-6 Alternate client and split mirror: primary client and alternate client share data through mirroring.



The figure shows the following phases:

Phase	Action
Phase 1	Mirror disk is synchronized with primary.


Phase	Action
Phase 2	Primary client collaborates with the alternate client to create the snapshot. Primary client splits the mirror disk from primary disk, and mounts the snapshot on the alternate client.
Phase 3	Alternate client streams the snapshot-based backup from the mirror to the media server.
Phase 4	Media server reads the backup image from the alternate client.

Figure 1-7 shows the media server and alternate client on the same host.

Figure 1-7 Alternate client and media server on same host

☒ Perform off-host backup
 Use: Alternate Client
 Machine: <alt_client_host>

(Detail from policy attributes dialog.)



The figure shows the following phases:

Phase	Action
Phase 1	Mirror disk is synchronized with primary.

Phase	Action
Phase 2	Primary client collaborates with the alternate client to create the snapshot. Primary client splits the mirror disk from primary disk, and mounts the snapshot on the alternate client.
Phase 3	Media server (serving as alternate client) reads the snapshot-based backup from the mirror.

A single alternate client can handle backups for a number of primary clients, as shown in the following diagram.

Multiple clients can share an alternate backup client of the same operating system type.

Note: All clients must be of the same OS.

Figure 1-8 Alternate client for multiple primary clients





Figure 1-9


Multiple clients with SSO: alternate client performs backup for multiple primary clients with NetBackup SSO option on a SAN



Alternate client backup through data replication example (UNIX only)

A volume that is configured with a software mirror on the alternate client replicates a volume that is on the primary client. When the backup starts, the replication is suspended. The software mirror is split from the replicating volume to form a snapshot on the alternate client. The snapshot is mounted on the alternate client and is used to complete the backup. After the backup, the snapshot volume is unmounted. The mirror is resynchronized with the replicating volume, and the replication is resumed.

Figure 1-10 Replication: primary client and alternate client share data through replication



Phase	Action
Phase 1	Primary client collaborates with the alternate client to create snapshot.
Phase 2	Alternate client creates the snapshot by splitting the mirror from a replicating volume.
Phase 3	Alternate client sends the snapshot data from the snapshot to the media server.
Phase 4	Media server reads the snapshot data from the alternate client.
Phase 5	Media server writes data to storage.

Only the VVR snapshot method for UNIX clients supports this configuration. This configuration requires the Veritas Volume Manager (VxVM version 3.2 or later) with the VVR license.

NetBackup media server data mover example (UNIX only)


In this off-host backup method, a NetBackup media server reads the backup data from the client snapshot and writes the data to a storage device. The NetBackup media server uses the mapping information that the client provides to read the data

as raw disk extents (units consisting of physical address and length). This method is not supported for Windows clients.

This mapping-based data movement is in addition to the normal backup processing that the media server performs in a master server or media server configuration.

Note: For a multi-ported SCSI disk array, a Fibre Channel SAN is not required.

Figure 1-11 NetBackup media server data mover



The figure shows the following phases in the backup process:

Phase	Action
Phase 1	On LAN, NetBackup master server tells the client to map the snapshot data on the disk.
Phase 2	On LAN, client sends the mapping information to the media server.
Phase 3	Media server processes the mapping information and reads client data over the SAN, from the addresses that the client provides.

Phase	Action
Phase 4	Media server writes data across the SAN to storage.

Snapshot Client requirements

NetBackup Snapshot Client requires the following components:

- A master server with NetBackup Snapshot Client server software installed.
- Clients running Solaris, HP, AIX, Linux, or Windows, with NetBackup Snapshot Client software installed.
Certain operating system and device patches (such as for the host bus adapter) may be required for both servers and clients.
See “[Snapshot Client assistance](#)” on page 32.

Please note the following additional requirements:

- For the disk array snapshot methods, assistance may be required from the disk array vendor.
- To use the snapshot and off-host backup features of NetBackup Snapshot Client with a NetBackup Oracle policy, UNIX clients must have Oracle8i or later installed.

Snapshot Client restrictions

For detailed support information, refer to the following:

- For a complete list of supported platforms, snapshot methods, data types, and database agents, and supported combinations of platform and snapshot methods, see NetBackup 7.x Snapshot Client Compatibility:
<http://www.netbackup.com/compatibility>
- Further information is also available.
See “[Snapshot Client assistance](#)” on page 32.

Note the following restrictions:

- Snapshot Client does not support the ALL_LOCAL_DRIVES entry in the policy’s **Backup Selections** list.
- For the NetBackup media server or Third-Party Copy Device methods, the client disk must be either a SCSI or Fibre Channel device.
- The NetBackup media server off-host backup method does not support the clients that use client deduplication. If the client is enabled for deduplication,

you must select **Disable client-side deduplication** on the policy **Attributes** tab.

Snapshot Client terminology

[Table 1-4](#) describes the terms that are used with NetBackup Snapshot Client. For explanations of other NetBackup terms, consult the NetBackup online glossary.

Table 1-4 Snapshot Client terminology

Term	Definition
SMDC	Snapshot Client
Alternate client backup	The alternate client performs a backup on behalf of another client.
Backup agent (see also Third-Party Copy Device)	A general term for the host that manages the backup on behalf of the NetBackup client. The agent is either another client, the NetBackup media server, a third-party copy device, or a NAS filer.
BCV	The mirror disk in an EMC primary-mirror array configuration (see mirror). BCV stands for Business Continuance Volume.
Bridge	In a SAN network, a bridge connects SCSI devices to Fibre Channel. A third-party copy device can be implemented as part of a bridge or as part of other devices. Note that not all bridges function as third-party copy devices.
Cache	Copy-on-write snapshot methods need a separate working area on disk during the lifetime of the snapshot. This area is called a cache. The snapshot method uses the cache to store a copy of the client's data blocks that are about to change because of file system activity. This cache must be a raw disk partition that does not contain valuable information: when you use the cache, the snapshot method overwrites any data currently stored there. See "How copy-on-write works" on page 215.
Copy Manager	See Third-Party Copy Device.
Copy-on-write	In NetBackup Snapshot Client, one of two types of supported snapshots (see also mirror). Unlike a mirror, a copy-on-write does not create a separate copy of the client's data. It creates a block-by-block "account" from the instant the copy-on-write was activated. The account describes which blocks in the client data have changed and which have not. The backup application uses this account to create the backup copy. Other terms and trade names sometimes used for copy-on-write snapshots are space-optimized snapshots, space-efficient snapshots, and checkpoints.
Data movement	A copy operation as performed by a third-party copy device or NetBackup media server.

Table 1-4 Snapshot Client terminology (*continued*)

Term	Definition
Data mover	The host or entity that manages the backup on behalf of the NetBackup client. The data mover can be either the NetBackup media server, a third-party copy device, or a NAS filer.
device	A general term for any of the following: LUN, logical volume, vdisk, and BCV or STD.
Disk group	A configuration of disks to create a primary-mirror association, using commands unique to the disks' vendor. See mirror and volume group.
Extent	A contiguous set of disk blocks that are allocated for a file and represented by three values: ■ Device identifier ■ Starting block address (offset in the device) ■ Length (number of contiguous blocks) The mapping methods in Snapshot Client determine the list of extents and send the list to the backup agent.
Fibre Channel	A type of high-speed network that is composed of either optical or of copper cable and employing the Fibre Channel protocol. NetBackup Snapshot Client supports both arbitrated loop and switched fabric (switched Fibre Channel) environments.
File system	
Instant Recovery	A restore feature of a disk snapshot of a client file system or volume. Client data can be rapidly restored from the snapshot, even after a system restart.
Mapping	Converting a file or raw device (in the file system or Volume Manager) to physical addresses or extents for backup agents on the network. NetBackup Snapshot Client uses the VxMS library to perform file mapping.
Mapping methods	A set of routines for converting logical file addresses to physical disk addresses or extents. NetBackup Snapshot Client includes support for file-mapping and volume-mapping methods.
Mirror	Has two meanings. ■ A disk that maintains an exact copy or duplicate of another disk. A mirror disk is often called a secondary, and the source disk is called the primary. All writes to the primary disk are also made to the mirror disk. ■ A type of snapshot that is captured on a mirror disk. At an appropriate moment, all further writes to the primary disk are held back from the mirror, which "splits" the mirror from the primary. As a result of the split, the mirror becomes a snapshot of the primary. The snapshot can then be backed up.

Table 1-4 Snapshot Client terminology (*continued*)

Term	Definition
NetBackup media server method	An off-host backup method in which the NetBackup media server performs the data movement.
Off-host backup	The off-loading of backup processing to a separate backup agent executing on another host. NetBackup Snapshot Client provides the following off-host backup options: Alternate Client, NetBackup media server, Third-Party Copy Device, and Network Attached Storage.
Primary disk	In a primary-mirror configuration, client applications read and write their data on the primary disk. An exact duplicate of the primary disk is the mirror.
Raw partition	A single section of a raw physical disk device occupying a range of disk sectors. The raw partition does not have a file system or other hierarchical organization scheme (thus, a "raw" stream of disk sectors). On some operating systems, such as Solaris and HP-UX, a raw partition is different from a block device over which the file system is mounted.
Recovery Manager (RMAN)	Oracle's backup and recovery program. RMAN performs backup and restore by making requests to a NetBackup shared library.
RMAN Proxy Copy	An extension to the Oracle8i media management API which enables media management software such as NetBackup to perform data transfer directly.
SAN (storage area network)	A Fibre Channel-based network connecting servers and storage devices. The storage devices are not attached to servers but to the network itself, and are visible to all servers on the network.
Secondary disk	See mirror.
Third-Party Copy Device	Has two meanings: ■ A backup agent on the SAN that operates on behalf of backup applications. The third-party copy device receives backup data from a disk that is attached to Fibre Channel and sends it to a storage device. The third-party copy device uses the SCSI Extended Copy command. The third-party copy device is sometimes called a Copy Manager, third-party copy engine, or data mover. In SAN hardware configurations, a third-party copy device can be implemented as part of a bridge, router, or storage device. The third-party copy device may or may not be the device to which the storage units are connected. ■ An off-host backup method in NetBackup Snapshot Client that allows backups to be made by means of a backup agent on the SAN.
VxMS (Veritas Federated Mapping Services)	A library of routines (methods) used by NetBackup Snapshot Client to obtain the physical addresses of logical disk objects such as files and volumes.

Table 1-4 Snapshot Client terminology (*continued*)

Term	Definition
Volume	A virtual device that is configured over raw physical disk devices (not to be confused with a NetBackup Media and Device Management volume). Consists of a block and a character device. If a snapshot source exists over a volume, NetBackup automatically uses a volume mapping method to map the volume to physical device addresses.
Volume group	A logical grouping of disks, created with the Veritas Volume Manager, to allow more efficient use of disk space.

Snapshot Client assistance

The following kinds of assistance are available.

Table 1-5 Types of assistance available for Snapshot Client

Type of assistance	Description
Snapshot Client help from NetBackup Administration Console	For help creating a policy, click the Master Server name at the top of the left pane and click Create a Snapshot Backup Policy .
Snapshot Client assistance from the web	For a document containing additional Snapshot Client assistance, see the tech note <i>NetBackup Snapshot Client Configuration</i>. This document may be accessed from the following link: http://www.veritas.com/docs/000081320 This document includes the following: ■ An up-to-date list of supported operating systems and peripherals ■ Detailed configuration assistance for the legacy disk array snapshot methods: EMC TimeFinder, Hitachi ShadowImage, and HP BusinessCopy ■ Sections on SAN device configuration and on setting up NetBackup for off-host data mover backups (with instructions for creating <code>3pc.conf</code> and <code>mover.conf</code> files)
Compatibility list	For a complete list of supported platforms, snapshot methods, data types, and database agents, and supported combinations of platform and snapshot methods, see the NetBackup 7.x Snapshot Client Compatibility document: http://www.netbackup.com/compatibility

About open file backups for Windows

The Open File Backup license is included in the standard NetBackup for Windows product, and enables open file backups of Windows clients. Open File Backup is independent of Snapshot Client. The Snapshot Client product is not required for Windows open file backups.

See the [NetBackup Administrator's Guide, Volume I](#).

Installation

This chapter includes the following topics:

- [Installation prerequisites for Snapshot Client](#)
- [Snapshot Client installation notes](#)
- [About distributing Client Software in mixed-platform environments](#)
- [About creating log directories](#)
- [About the snapshot state file](#)

Installation prerequisites for Snapshot Client

Note the following prerequisites:

- NetBackup Enterprise server 10.0 or later must be installed on the master and the media servers. For performing backups of the primary client, the master or the media server can be running any supported UNIX, Linux or Windows platform. For a list of supported platforms for Snapshot Client, refer to the following document:
<http://www.veritas.com/docs/000081320>
- From NetBackup version 9.1 onwards, the media server should be running the same NetBackup version as that of the master server for it to communicate with the CloudPoint server
<http://www.veritas.com/docs/000035749>
- NetBackup 10.0 or later client software must be installed on clients.

Snapshot Client installation notes

The following list contains C installation information for RHEL, UNIX and Windows.

RHEL

Perform the following tasks for RHEL:

- In the recent RHEL version, there are changes in usage of partition delimiter. Based on the last character in the device path, OS decides if partition delimiter must be added or not. To address the change, update the udev rule to add partition delimiter 'p' on RHEL.

- Open the `/lib/udev/rules.d/62-multipath.rules` file.

- Locate these line:

```
RUN+="$env{MPATH_SBIN_PATH}/kpartx -an $tempnod",
\ ENV{DM_MULTIPATH_NEED_KPARTX}=""
```

Update the lines as follows:

```
RUN+="$env{MPATH_SBIN_PATH}/kpartx -an -p p $tempnod",
\ ENV{DM_MULTIPATH_NEED_KPARTX}=""
```

- Restart the multipath service, using the following command:

```
# service multipathd restart
```

- For devices in LVM, if multiple paths are used, ensure that the multipathing service is enabled.
- Junction mount point on filler or array is not supported for snapshot operations with VSO FIM snapshot method.
- For client-side nested mount points, rollback of parent mount points is not supported.

UNIX

- NetBackup Snapshot Client is installed with the NetBackup client software. Every NetBackup server includes the NetBackup client software, by default. So, you can use NetBackup Snapshot Client on a NetBackup server or client, if the Snapshot Client is supported on that platform.
- For NetBackup 10.0, Snapshot Client for Solaris is supported on SPARC computers only.
 For the NetBackup installation procedure, refer to the [NetBackup Installation Guide](#)
- If you install in a cluster environment, first freeze the active node so that migrations do not occur during installation.
 For information about freezing a service group, see the clustering section in the [NetBackup High Availability Administrator's Guide](#) for the cluster software you are running.

- The package provides the `rescan-scsi-bus.sh` script which helps to rescan and discover new devices after a snapshot is imported. This script is used during 'bpfis import' and 'bpfis deport' operations. If this script is not present on the NetBackup client, snapshot import fails.

The default location of the script is:

```
<install_location>/usr/bin/rescan-scsi-bus.sh
```

In case, the package is not deployed run the following command:

```
yum install sg3_utils
```

- For Linux LVM, ensure that `systemctl` is available.
- Ensure that F user is configured.

Refer to the uninstall procedure that is described in the [NetBackup Installation Guide](#)

About distributing Client Software in mixed-platform environments

For UNIX and Windows, Snapshot Client software is automatically installed with the base NetBackup client software. See the appropriate [NetBackup Installation Guide](#) for more information.

For mixed environments, note the following:

- If you have a Windows server with UNIX clients, you must install the client software on the UNIX client computers individually, from the NetBackup media.

About creating log directories

During backup and restore, Snapshot Client messages are written to several log directories on the NetBackup server and client, if the directories exist. For logging to occur, you must create these directories manually if the directories do not already exist.

See [“About gathering information and checking logs”](#) on page 177.

About the snapshot state file

Whenever NetBackup creates a snapshot, it also creates a file that contains information about the snapshot. This file is called the snapshot state file. The state file stores information that is required for performing certain operations on the snapshot. (Examples of such operations are, data restore from snapshots, backup from snapshots, or snapshot deletion.

The NetBackup bpfis process creates a state file on the client. A copy of the state file is stored on the NetBackup primary server. In a cluster, if the client experiences problems and failover to another node occurs, the state file on the NetBackup primary server comes in handy to obtain information about the failed client. In case the state file does not exist on the NetBackup primary server, it becomes impossible for the active node to get snapshot information.

Policy configuration

This chapter includes the following topics:

- [Notes on Snapshot Client policies](#)
- [Configuring a Snapshot Client policy](#)
- [Backup Selections tab options when configuring a policy](#)
- [Off-host backup configuration options](#)
- [Selecting the snapshot method](#)
- [Configuring backup scripts](#)
- [About using alternate client backup](#)
- [Configuring alternate client backup](#)
- [Policy configuration tips](#)
- [About disabling snapshots](#)

Notes on Snapshot Client policies

Before you configure a Snapshot Client policy, note the following.

- NetBackup product must be installed on primary server(s) and clients.
- For the NetBackup Media Server off-host backup methods, a Fibre Channel network or multi-ported SCSI disk array must be configured.
- Storage devices must be configured (you can use the Device Configuration Wizard).

- Encryption and compression are supported, but are applied only to the backup copy that is written to a storage unit. The snapshot itself is neither compressed nor encrypted.
- Ensure that the number of LUNs exposed to the HBAs, do not reach the maximum limit during any snapshot related operations.

Configuring a Snapshot Client policy

The following procedure describes only the options pertaining to a Snapshot Client policy.

To configure a Snapshot Client policy

- 1 On the primary server, open the NetBackup Administration Console.
- 2 Click **Policies** and select the policy to edit.

Change Policy - Snap_policy1

Server: nbu-id1

Attributes Schedules Clients Backup Selections

Policy type: MS-Windows

Destination: <No data classification>

Data classification: <No data classification>

Policy storage: Any_available

Policy volume pool: NetBackup

☐ Take checkpoints every: 0 minutes

☐ Limit jobs per policy:

Job priority: 0 (higher number is greater priority)

Media Owner: Any

Snapshot Client and Replication Director

☐ Perform block level incremental backups

☐ Use Replication Director

☒ Perform snapshot backups [Options...](#)

☐ Retain snapshot for Instant Recovery or SLP management

☐ Hyper-V server:

☐ Perform off-host backup

Use:

Machine:

☒ Go into effect at: 02/23/2015 18:34:17

☐ Backup network drives

☐ Cross mount points

☐ Compress

☐ Encrypt

Collect disaster recovery information for:

☐ Bare Metal Restore

☐ Collect true image restore information

☐ with move detection
(Required for synthetic backups and Bare Metal Restore)

☐ Allow multiple data streams

☐ Disable client-side deduplication

☐ Enable granular recovery

☐ Use Accelerator

☐ Enable optimized backup of Windows deduplicated volumes

Keyword phrase (optional):

Microsoft Exchange Server Attributes

Exchange DAG or Exchange 2007 replication (LCR/CCR)

Database backup source:

Preferred server list... (Exchange DAG only)

OK Cancel Help

- 3 Select the policy type:
 - NetBackup 10.0 supports off-host backup of Oracle database in the SFRAC (Storage Foundation Real Application Clusters) environment. For more details, refer to the *NetBackup for Oracle Administrator's Guide*:

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

4

- 5 Make sure **Perform snapshot backups** is selected.

Note: When you select **Perform snapshot backups**, the **Bare Metal Restore** option is disabled.

Note: **Perform snapshot backups** must be selected for the policy to reference any storage lifecycle policy with a Snapshot destination.

- 6 Select the snapshot method manually.

See [“Selecting the snapshot method”](#) on page 43.

- 7 Select the **Retain snapshots for instant recovery or SLP management** attribute.

This attribute is required for block-level restore, file promotion, and rollback.

See [“Instant Recovery restore features”](#) on page 156.

See [“Configuring a policy for Instant Recovery”](#) on page 87.

- 8 To reduce the processing load on the client, select **Perform off-host backup**.

See [“Off-host backup configuration options”](#) on page 42.

- 9 To save these settings, click **Apply**.

- 10 To define a schedule, use the **Schedules** tab, and to specify the clients, use the **Clients** tab .

Regarding clients: only one snapshot method can be configured per policy. To select one snapshot method for clients a, b, and c, and a different method for clients d, e, and f: create a separate policy for each group of clients and select one method per policy. You may be able to avoid this restriction using the auto method.

- 11 To specify the files to be backed up, use the **Backup Selections** tab .

See [“Backup Selections tab options when configuring a policy”](#) on page 42.

- 12 On the **Policy Attributes** tab: if you click **Apply** or **OK**, a validation process checks the policy and reports any errors. If you click **Close**, no validation is performed.

Backup Selections tab options when configuring a policy

Note the following about the options on the **Backup Selections** tab:

- Snapshot Client policies do not support the ALL_LOCAL_DRIVES entry in the **Backup Selections** list (except for the VMware and Hyper-V snapshot methods).
- For snapshots, the maximum pathname length is approximately 1000 characters (as opposed to 1023 characters for backups that do not use a snapshot method). See [“Maximum pathname length”](#) on page 49. The [NetBackup Administrator’s Guide, Volume I](#), describes other file-path rules.
- Wildcards are permitted if the wildcard does not correspond to a mount point or a mount point does not follow the wildcard in the path.

Note: This is applicable to a Storage Lifecycle Policy that has snapshot as the first operation and does not contain any backup or replicate operation.

For example, in the path `/a/b`, if `/a` is a mounted file system or volume, and `/a/b` designates a subdirectory in that file system: the entry `/a/b/*.pdf` causes NetBackup to make a snapshot of the `/a` file system and to back up all pdf files in the `/a/b` directory. But, with an entry of `/*` or `/*/b`, the backup may fail or have unpredictable results, because the wildcard corresponds to the mount point `/a`. Do not use a wildcard to represent all or part of a mount point.

In another example, `/a` is a mounted file system which contains another mounted file system at `/a/b/c` (where `c` designates a second mount point). A Backup Selections entry of `/a/*/c` may fail or have unpredictable results, because a mount point follows the wildcard in the path.

Information is available on the **Cross mount points** policy attribute.

See [“Snapshot tips”](#) on page 49.

Off-host backup configuration options

Off-host backup shifts the burden of backup processing to a separate backup agent, which reduces the impact on the client’s resources ordinarily caused by a local backup. The client supplies a small amount of information on how data is mapped. The backup agent does the bulk of the work by sending the client’s data to the storage device.

Select the off-host backup method when you create a new policy, you can also select off-host backup for an existing policy. Select off-host backup from the Add New Policy dialog box. The following image shows a part of the dialog box.

See “Configuring a Snapshot Client policy” on page 39.

Snapshot Client and Replication Director

☐ Perform block level incremental backups

☐ Use Replication Director

☒ Perform snapshot backups Options...

☐ Retain snapshot for Instant Recovery or SLP management

☐ Hyper-V server:

☒ Perform off-host backup

Use:

Machine:

Alternate client
Alternate client
Data mover
VMware backup host

The **Use** and **Machine** fields designate the backup agent.

Select the type of off-host backup from the **Use** field:

- **Alternate Client**
- **Data Mover**

Select this option to designate the backup agent as a NetBackup media server, a third-party copy device that implements the SCSI Extended Copy command, or a NAS filer (Network Attached Storage).

The **Data Mover** option requires the Standard, FlashBackup, or MS-Windows policy type.

Select the type of data mover in the **Machine** drop-down:

Selecting the snapshot method

You can manually select a snapshot method as follows.

To select a snapshot method

- 1 On the primary server, open the NetBackup Administration Console.
- 2 Click **Policies**.
- 3 In the **All Policies** pane, double-click the name of the policy.
The **Change Policy** dialog appears.
- 4 Make sure **Perform snapshot backups** is selected.
- 5 Click **Options**.

The following selections are optional. If no snapshot method is selected, NetBackup will select a snapshot method at the time of backup.

Snapshot method for this policy:
FlashSnap

The following parameters can be set for backups performed in this policy that use this snapshot method:

Configuration Parameters

Parameter	Value
Keep snapshot after backup	No
Resynchronize mirror in background (not Instant Recovery)	No
Maximum Snapshots (Instant Recovery only)	1

Snapshot Resources

Array SN	Source	Snapshot Devices
----------	--------	------------------

Add Change Remove Remove All

OK Cancel Help

- 6 In the pull-down menu, select **VSO** as the **Snapshot method** for the policy, to manage snapshots using the CloudPoint.

Only one snapshot method can be configured per policy.

- 7 Specify required parameters, if any. Additional fields in the **Snapshot Options** dialog indicate required parameters.

See [“Configuration parameters for Snapshot Client”](#) on page 45.

See [“Snapshot Resources”](#) on page 46.

- 8 Click **OK**.

Configuration parameters for Snapshot Client

The following parameters are required for certain snapshot methods. If no additional parameters are required for the snapshot method you have chosen, the **Snapshot Options** dialog states that no parameters are required.

Maximum Snapshots parameter (Instant Recovery only)

One of several options for determining when Instant Recovery snapshots are deleted.

Policy validation fails if there is a mismatch of retention found on the snapshot. For example, if the **Maximum Snapshots (Instant Recovery only)** parameter is set to any value other than Managed by SLP and the SLP used in the same policy has **Fixed** retention for the Snapshot job the policy validation fails. If you have such a policy configured on a pre- 7.6 NetBackup primary server, it is advisable that you validate and correct the policy after you upgrade to a NetBackup 10.0 primary server.

See [“Means of controlling snapshots”](#) on page 87.

Important note on VSS and disk arrays

To back up a Windows client with the VSS method, please note the following about snapshot parameter values:

- If you selected a **Provider Type** of **3 (hardware)** and a **Snapshot Attribute** of **2 (plex)**: You must configure an appropriate number of clones or mirrors in the disk array, depending on the value that is specified for the **Maximum Snapshots (Instant Recovery only)** parameter. If the **Maximum Snapshots** parameter is set to **3**, you must configure three clones or mirrors. If you do not want to use Instant Recovery and the **Maximum Snapshots** parameter is **1**, you need only configure one clone or mirror.
- You must also synchronize the clones or mirrors with the disk array source device before starting the backup. If the clones or mirrors are not synchronized before the backup begins, VSS cannot select a clone or mirror on which to create the snapshot. The backup fails.

Snapshot Resources

To configure the disk array methods, see the chapter titled Configuration of snapshot methods for disk arrays:

Configuring backup scripts

For backups using a snapshot method, you can run scripts before and after the snapshot by adding directives to the **Backup Selections** list, as follows.

To configure backup scripts

- 1 On the primary server, open the NetBackup Administration Console.
- 2 Click **Policies**.
- 3 In the **All Policies** pane, double-click the name of the policy.
The **Change Policy** dialog appears.
- 4 Click the **Backup Selections** tab.
- 5 Add the following directive to the start of the **Backup Selections** list:

```
METHOD=USER_DEFINED
```


- 6 Add one or both of the following directive(s), as required.

```
DB_BEGIN_BACKUP_CMD=your_begin_script_path  
DB_END_BACKUP_CMD=your_end_script_path
```

Note: DB_BEGIN_BACKUP_CMD and DB_END_BACKUP_CMD directives are not supported for Windows clients.

The arguments (arg) are optional.

For example:



In this example, the script `shutdown_db.ksh` is run before the backup, and `restart_db.ksh` is run after the snapshot is created.

About using alternate client backup

Alternate client backup is off-host in that all backup processing is off-loaded to another client. Off-loading the work to another client saves computing resources on the original client. The alternate client handles the backup I/O processing and the backup has little effect on the original client.

Alternate client backup requirements

Before you configure a policy for alternate client backup, make sure the following have been done:

- The client data must be configured as required by the snapshot method. The alternate client must have access to the snapshot devices.
- User and group identification numbers (UIDs and GIDs) for the files to back up must be available to the primary client and the alternate client.
- The primary client and alternate client must run the same version of NetBackup. For example, the use of a later version of NetBackup on the primary client and an earlier version on the alternate client is not supported.

- The primary client and alternate client must run the same operating system, volume manager, and file system. For each of these I/O system components, the alternate client must be at the same level as the primary client, or higher level.

Table 3-1 lists the supported configurations.

Table 3-1 Alternate client requirements

If primary client is:	Alternate client must be:
Solaris	Solaris, at same level as primary client or higher
HP	HP, at same level as primary client or higher
AIX	AIX, at same level as primary client or higher
Red Hat	Red Hat, at same level as primary client or higher
SUSE	SUSE, at same level as primary client or higher

Configuring alternate client backup

For an alternate client backup policy, make the following selections. Some of the options depend on your hardware configuration and product licenses.

To configure a policy for alternate client backup

- 1 For **Policy type**, choose **Standard**, **FlashBackup**, **FlashBackup-Windows**, **MS-Windows**, **MS-Exchange-Server**, **MS-SQL-Server**, **DB2**, **SAP**, or **Oracle**.
- 2 Click **Perform snapshot backups**.
- 3 Click **Perform off-host backup**.
- 4 Click **Use alternate client** and select the alternate client from the drop-down list or type it in.
- 5 Click **Options** if you want to choose the snapshot method.

Snapshot method: You can select the auto method, or the following:

- FlashSnap, for a disk group split configuration, with VxVM 3.2 or later with the FlashSnap feature.
- VVR, for a UNIX replication host; requires VxVM 3.2 or later with VVR feature.
- VSS, for snapshots using the Volume Shadow Copy Service of Windows. This method is for Windows clients, where the client data is stored on a disk array such as EMC or Hitachi, or in a Veritas Storage Foundation for Windows 4.1 or later volume. Supports Exchange.

- The disk array-related snapshot methods.

Example alternate client backup configurations

Consider the following example configurations:

Table 3-2 Alternate client backup configurations

Configuration	Description
Client data is on a Dell EMC disk array	If the data is in an Oracle database, select Oracle as the policy type.

Policy configuration tips

The following topics relate to policy creation.

Maximum pathname length

For snapshots, the maximum file list pathname length is approximately 1000 characters (1023 characters for the backups that do not use a snapshot method). The reason is that the snapshot is created on a new mount point which is derived from the original **Backup Selections** list path name. If the new mount point plus the original file path exceeds the system-defined maximum path name length (1023 characters on many systems): the backup fails with a status code 1, "the requested operation was partially successful."

Refer to the [NetBackup Administrator's Guide, Volume I](#), for other NetBackup file-path rules.

Snapshot tips

Note the following tips:

- In the **Backup Selections** list, be sure to specify absolute path names. Refer to the [NetBackup Administrator's Guide, Volume I](#) for help specifying files in the **Backup Selections** list.
- If an entry in the **Backup Selections** list is a symbolic (soft) link to another file, Snapshot Client backs up the link, not the file to which the link points. This NetBackup behavior is standard. To back up the actual data, include the file path to the actual data.
- The **Cross mount points** policy attribute is not available for the policies that are configured for snapshots. This option is not available because NetBackup

does not cross file system boundaries during a backup of a snapshot. A backup of a high-level file system, such as / (root), does not back up the files residing in lower-level file systems. Files in the lower-level file systems are backed up if they are specified as separate entries in the **Backup Selections** list. For instance, to back up /usr and /var, both /usr and /var must be included as separate entries in the **Backup Selections** list.

For more information on **Cross mount points**, refer to the [NetBackup Administrator's Guide, Volume I](#).

Multiple data streams

Make sure that these settings allow the number of active streams to be equal to or greater than the number of streams in the **Backup Selections** list:

- Policy attribute
Limit jobs per policy
- Schedule setting
Media multiplexing
- Storage unit setting
Maximum multiplexing per drive
- System configuration setting
Maximum jobs per client

About disabling snapshots

This section explains how to turn off the creation of snapshots for certain clients but to continue making non-snapshot backups of those clients.

Disabling Snapshot Client snapshots

Use the following procedure to disable Snapshot Client snapshots.

Warning: All features of Snapshot Client depend on snapshots. Clearing the **Perform snapshot backups** check box disables all features of Snapshot Client, such as off-host backup and Instant Recovery. Backups that depend on these features are disabled.

To disable Snapshot Client snapshots for individual clients

- 1** Determine which Snapshot Client policies are configured for the clients on which you want to turn off snapshots.
- 2** In the NetBackup Administration Console, click **NetBackup Management > Policies**.
- 3** In the **All Policies** pane, double-click a Snapshot Client policy.
- 4** In the **Snapshot Manager** pane of the **Attributes** tab, clear the check box for **Perform block level incremental backups**.
- 5** In the same pane, clear the check box for **Perform snapshot backups**.

NAS snapshot configuration

This chapter includes the following topics:

- [About NAS snapshot overview](#)
- [Dynamic data streaming for D-NAS workloads](#)
- [Prerequisites for NetBackup DNAS configuration](#)
- [Configuring dynamic data streaming with backup host pool](#)
- [Configuring a backup host pool](#)
- [Discover plugins](#)
- [Minimum supported backup host versions for different features](#)
- [Limitations and considerations](#)
- [Setting up a NAS-Data-Protection policy](#)

About NAS snapshot overview

- **Dynamic NAS (D-NAS):** By means of Snapshot management server and the storage array plugins, NetBackup can make snapshots of NAS volumes and shares. The snapshots are accessed on backup hosts and read by dynamic streams to create point in time backup copies.
- **NAS snapshot through staged NAS volume or share:** NetBackup can make snapshots of client data that resides on the NAS host and be accessible on the client through an NFS mount on UNIX or CIFS on Windows.

A NAS snapshot is a point-in-time disk image. Snapshots can be retained on disk as long as desired. The data can be efficiently restored from disk by means of the Snapshot Client Instant Recovery feature.

In the NetBackup policy, enter the following:

- For Windows client:
- For UNIX client:

`/NFS_mountpoint`

Note: Windows pathnames must use the Universal Naming Convention (UNC).

NetBackup creates snapshots on the NAS-attached disk only, not on the storage devices that are attached to the NetBackup server or the client.

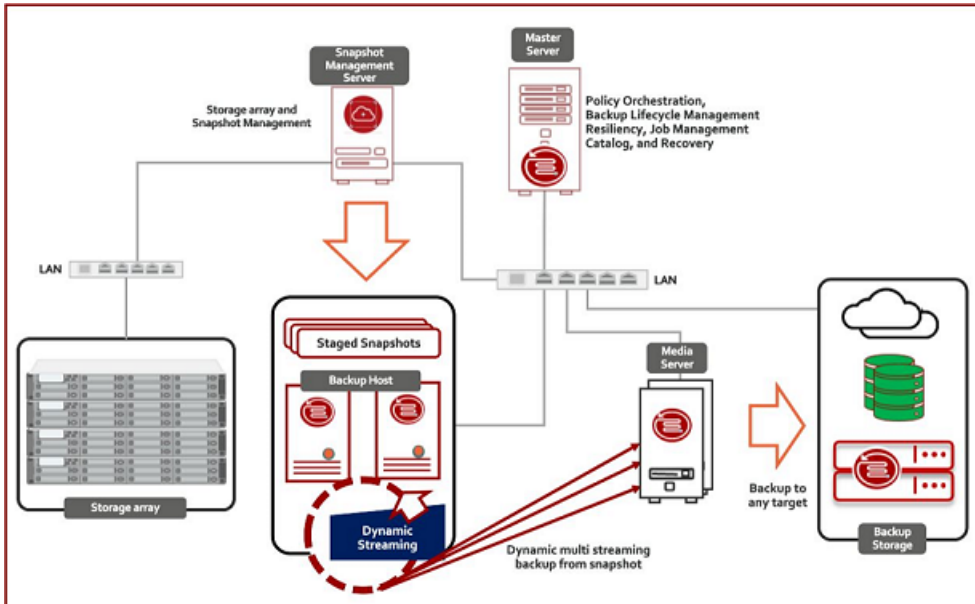
Dynamic data streaming for D-NAS workloads

You can perform a snapshot enabled, off-host backup of NAS volumes, where a volume is backed up using dynamic backup streams.

Each NAS volume or share is read over NFS or SMB, and backed up using a configured number of backup streams. Files of these NAS volumes or shares are optimally distributed real time across streams to realize the full potential of backup streams. You cannot mix NAS volumes or shares of different storage array vendors in a single policy. In other words, using a single policy you can only protect assets for a single vendor and single NAS protocol.

Dynamic streaming is built on the NetBackup client framework and uses NAS-Data-Protection policy type for snapshot and backup orchestration of NAS data. This policy supports SLP only for data lifecycle.

Figure 4-1 Dynamic data streaming



Understanding the features

The following table explains the key features:

Table 4-1 Feature details

Term	Description
Backup host pool	Backup host pool is a group of NetBackup backup hosts where the snapshot of the volume is staged for backup process to read. These hosts can be NetBackup client, media, or primary server.

Table 4-1 Feature details *(continued)*

Term	Description
Vendor change tracking	Vendor change tracking (VCT) is a mechanism to get the difference in the content of the volume or share between two point in time snapshots. It relies on storage array vendor's native technology to identify the difference, that is, add, modify, and delete files between two point in time snapshots of the same volume. You must ensure that the storage array you are using provides such a capability in order to use this feature. VCT is not applicable in the following conditions: ■ Schedule type is other than INCR and CINR. It's only supported for INCR and CINR. ■ Base snapshot is not available. ■ Expire after copy retention options is selected for snapshot in SLP. ■ Accelerator is enabled for the policy.
Exclude volumes	The volumes out of the backup selection list that you do not want to backup. For example, /prodVol* is the backup selection and there is a volume /prodVol-Scratch which you do not want to backup.
Accelerator	NetBackup's traditional Accelerator feature can be leveraged along with dynamic streaming for optimized and fast backups.
Checkpoint Restart	You can leverage NetBackup's traditional Checkpoint restart feature along with dynamic streaming. By taking checkpoints periodically during the backup, NetBackup can retry a failed backup from the beginning of the last checkpoint without restarting the entire job.

Dynamic streaming parameters

Dynamic streaming as the name suggests is a group of backup streams running in parallel which dynamically distributes files for backups amongst them for optimized and fast backup of dense NAS volumes or shares.

- **Maximum number of streams per volume:** The value will determine the number of backup streams that will be deployed for backing up each volume.
 For example, If a policy contains 10 volumes and the value of this parameter is set to 4, then you will see group of 4 backup streams for each volume, thereby total of 40 child backup streams and 10 parent backup streams as part of backup execution of the policy.
- **Maximum number of files in a batch:** The value will determine the maximum number of files that will be processed in a single batch by any stream. The files in this batch will be processed sequentially within the stream.
 For example, value of 300 for this parameter will mean that every stream will be assigned maximum of 300 files in a single batch. So if a volume has one million files to be backed up and 4 streams assigned, then every stream will be assigned 300 files each to begin with and then subsequently the streams will be fed with 300 more files as and when they are ready for more backup processing.

Prerequisites for NetBackup DNAS configuration

Note the following prerequisites:

- Ensure that you have installed the CloudPoint server. For details refer to the *NetBackup CloudPoint Install and Upgrade Guide*.
- Prepare the plug-in that you want to use for the NetBackup DNAS configuration. See the *NetBackup CloudPoint Install and Upgrade* guide for additional SMB configuration.
- Identify the backup host that you want to use for the configuration.
- If NAS Data Protection policy uses TAPE storage unit in SLP for protecting NAS volumes, then the number of tape drives must be greater than or equal to the maximum number of streams per volume, otherwise backups fail. The other parameters of TAPE, like Media multiplexing and maximum concurrent write drives, does not have any affect on NetBackup DNAS backups.
- For SMB backups using NAS-Data-Protection policy the master, media and backup host version should be 9.1 onwards.

Logging on to the NetBackup services as a domain user

This step is required for Windows backup hosts for NetBackup DNAS-SMB only. You must log on to the NetBackup client service and the NetBackup legacy network service as a domain user to perform the tasks described in the following sections.

Note: The Windows domain user must be a part of the local Administrative group.

To log on to the NetBackup services as a domain user:

- 1 Make sure that the NetBackup client service and the NetBackup legacy network services are running.
- 2 In Windows Services, double-click the NetBackup service.
- 3 Check the **Log On** tab: if any of this service is not logged on as the domain user, change the logon to the domain account and restart the service. If both the services are not logged on as the domain user, you must do it in the following sequence:
 - Log on to the first service as domain user, and restart the service.
 - Log on to the second service as domain user, and restart the service.
- 4 Make sure that all NetBackup services are running.
- 5 Relaunch the NetBackup UI.

See [“About NAS snapshot overview”](#) on page 52.

Configuring dynamic data streaming with backup host pool

Backup host pool is a group of NetBackup backup hosts where the snapshot of the volume is staged for backup process to read. These hosts can be NetBackup client, media, or primary server. Although NetBackup client is not listed in the available backup hosts but it can be added explicitly to the list.

Note: Ensure that all hosts in a backup host pool must have same the OS, OS path level, and default NFS version configuration.

Table 4-2 Configuration steps and description

Step	Description	Reference
1	Register the CloudPoint server and add it as a snapshot management server in NetBackup.	
2	Configure the NAS plug-in.	For details refer to the <i>NetBackup CloudPoint Install and Upgrade Guide</i> .
3	Add the backup hosts to a backup host pool. The backup hosts are responsible for data streaming.	See “Configuring a backup host pool” on page 58.


Table 4-2 Configuration steps and description (*continued*)

Step	Description	Reference
4	Configure a storage lifecycle policy for snapshots and backup from snapshots also.	See “About storage lifecycle policies for snapshots” on page 90.
5	Configure the NAS-Data-Protection policy.	See “Setting up a NAS-Data-Protection policy” on page 64.

Configuring a backup host pool

Backup hosts and backup host pools are used for NAS-Data-Protection policy based on dynamic multistreams.

Figure 4-2 Backup host pool



You can use a NetBackup primary server, media server, or a standalone client as a backup host. For the hosts that you add to the backup host pool, their volumes are distributed for backup purposes on the backup hosts. This configuration results in a better backup performance.

Note: A NetBackup primary server running on Veritas Flex Appliance is not supported as a backup host for a NAS-Data-Protection policy.

You can create a backup host pool with different versions of NetBackup hosts. You can create Windows backup host pools only with version 9.0.1 or later. Windows hosts with a version earlier than 9.0.1 are not displayed.

Note the following important points:

- In a backup host pool you can either have Linux hosts or Windows hosts only. A pool does not support hosts with both platforms.
- All the hosts in the backup host pool must use the same OS version. This way each host has the same version of NFS for consistent backups.
- For backup hosts with a multi-NIC setup, add the host name that is already used on the NetBackup primary server. Do not add an alias name or any other host names in the backup host pool.

To configure a backup host pool

- 1 Open the **NetBackup Administration Console**.
- 2 Select **NetBackup Management > Host Properties > primary server**.
- 3 Double-click on the preferred primary server host name.
- 4 Click **Backup Host Pools**.
- 5 Click **Add**.
- 6 In the **Add Backup Host Pool** dialog box, enter a host pool name.
- 7 (Conditional) This step is applicable only for the clients that you want to add to the list. In the **Enter hostname to add to the list** field, add the client name and click **Add to list**.
- 8 Select the **OS Type**.
- 9 Select the backup hosts that you want to add to the list.
- 10 Click **OK**.

Note: You cannot delete a backup host pool, if it is configured with an existing NAS-Data-Protection policy.

Discover plugins

Discover plugin is a generic ISM feature that is available for all array plugin's whether it is a NAS or SAN plugin. This is not specific to NAS-Data-Protection policy.

You can trigger Plugin discovery from Java UI or using the `tpconfig` CLI. When you trigger plugin discovery, JAVA UI internally invokes `tpconfig` to perform the discovery operation. This operation is synchronous. It does not return till the discovery operation is complete.

To discover plugins:

- Get all the configured plugins from the `CloudPoint_plugin.conf` file.
- Get the unique plugin types from the configured plugins.
Get all configured plugins from CloudPoint for every plugin type (1 CloudPoint API for 1 plugin type).
- Update existing plugin info in-memory (discovery status and time).
- Persist the updated plugin info in `CloudPoint_plugin.conf` file.
- Return the updated plugin info to the User/JavaUI.

Run the command:

```
/usr/opensv/volmgr/bin/tpconfig -discover_plugin -cloudpoint_server  
<cloudpoint_server_name> -plugin_id <plugin_id_name>
```

```
# time tpconfig -discover_plugin -c "10.10.10.10" -plugin_id nut1  
  
Waiting for the plug-in discovery to complete...  
  
Discovery is complete for plug-in [nut1].  
  
real    0m32.429s  
user    0m0.078s  
sys     0m0.023s
```

Discovery status (CLI):

```
>>> time tpconfig -list_plugins -c "[redacted] server"
=====
Plugin details for 'V-L [redacted]'
=====
1. Plugin_ID      : netapp1
   Plugin_Type    : netapp
   Use_VCT        : false
   Discovered_Time : Mon Mar 22 10:55:27
   Discovery_Status : discovering
=====

real    0m0.384s
user    0m0.078s
sys     0m0.018s

>>> time tpconfig -list_plugins -c "[redacted] server"
=====
Plugin details for "[redacted]"
=====
1. Plugin_ID      : netapp1
   Plugin_Type    : netapp
   Use_VCT        : false
   Discovered_Time : Mon Mar 22 15:14:39
   Discovery_Status : discovered
=====

real    0m0.364s
user    0m0.082s
sys     0m0.015s

>>> time tpconfig -list_plugins -cloudpoint_server "[redacted]"
netapp1 netapp Array "NetApp Storage Array" 10.210.231.58 NAS,SAN 1616406279.338631 discovered

real    0m0.393s
user    0m0.083s
sys     0m0.020s
```

Performing discovery in Java UI

To perform discovery in Java UI

- 1 On the left, expand **Media and Device Management**> **Credentials** > click **Snapshot Management Servers**.
- 2 Right-click the required CloudPoint plugin, click **Discover assets**.

You can view the discovery status for different plugins in NetBackup Admin console.

Minimum supported backup host versions for different features

Different features of NAS Data Protection policy requires backup host with NetBackup version greater than or equal to the minimum supported backup host version. The following table specifies which feature is supported from which NetBackup version.

Table 4-3 NAS data protection policy features

Supported features	Minimum supported backup host version
Only NFS backup	8.3
NFS and Vendor change tracking	8.3
NFS and Checkpoint restart enabled backups	9.0
NFS and Accelerator enabled backups	9.1
NFS, Checkpoint restart, and Accelerator	9.1
SMB backups (including CPR, accelerator, Vendor change tracking)	9.1

Limitations and considerations

You can set up a NAS-Data-Protection policy for your workloads.

Note: If you use cloud as a storage unit, you must configure appropriate buffer size. Refer to the *NetBackup Cloud Administrator's Guide*.

Note the following important points about NAS-Data-Protection policy.

- The policy is supported for Nutanix, NetApp, Qumulo, Dell EMC Power Scale (Isilon), and Dell EMC Unity plug-ins only.
- The NAS-Data-Protection is not supported in the DNAT environment.
- This policy does not support copy-based retention for Snapshot images. Ensure that you carefully plan your policy scheduling and snapshot retention in SLP.
- Client side deduplication is not supported for NAS-Data-Protection policy.
- Vendor Change Tracking (VCT) and Accelerator options are mutually exclusive for NAS-Data-Protection policy. You cannot enable both. Veritas recommends

that you do not toggle the policy with these options between different execution of this policy.

- Vendor Change Tracking (VCT) enabled backup with incremental schedule requires base snapshot copy to determine the difference between current snapshot copy and base snapshot copy. Differential incremental schedule refers to base snapshot copy from previous differential incremental or cumulative incremental or full schedule. Cumulative incremental schedule refers to base snapshot copy from previous cumulative incremental or full schedule. During VCT enabled backup with incremental schedule, if the base snapshot copy is not available then the backup operation might fail with the error shown in Activity Monitor Detailed status.
- NAS-Data-Protection policy is a snapshot enabled data protection policy. You can configure only storage lifecycle policy (SLP) against policy's storage destination. Additionally, the SLP should always have Snapshot as the primary job and Backup from Snapshot as secondary job.
- If the NAS-Data-Protection policy is used in a backup host that is running anti-virus software, the parent backup from snapshot job might hang. The anti-virus software may block NetBackup process interactions causing the processes to hang. In this particular scenario, the nbcs process on the backup host might hang resulting in the backup-from-snapshot job to hang. Create an antivirus exclusion for nbcs on the backup host.
 To cancel the hung job:
 - Note down the process ID of the nbcs process which is running on the backup host. This can be obtained from the job details section.
 - Login to the backup host and manually kill the nbcs process.
 - Refer to the Technote for more details regarding how to exclude the NetBackup processes from virus scanning:
https://www.veritas.com/support/en_US/article.100004864
 - If the above steps cannot resolve the issue (and the nbcs hang persists), uninstall the network component from antivirus. On Symantec Endpoint Protection, this is called the "Network and Host Exploit Mitigation" component.
- For NAS-Data-Protection policy, multiple images are created for a single volume that is backed up. The number of images is equal to the value configured for the **Maximum number of streams per volume** in the policy. Since a single image cannot be referred from a single volume, NetBackup groups the images associated with a volume. When an operation is performed on one of the images in a volume, the same operation is also performed on the other grouped images in the volume. For example, if **Maximum number of streams per volume** is

set as four and you select one image for a volume to expire, the other three images also expire. The image grouping is applicable for the following operations:

- Browse and Restore
- Image expiration
- Image import
- Image duplication
- Image verification
- Set primary copy

Note: Image grouping is not applicable for importing images as part of Image Sharing operation.

- To enable checkpoint restart for NAS-Data-Protection policies created before upgrading to version 9.0, you must select the **Take checkpoints every** check box and enter a value in minutes.

Note: The NAS-Data-Protection policy cannot be configured using the NetBackup web UI.

Setting up a NAS-Data-Protection policy

To set up a policy for NAS data protection

- 1 Open the **NetBackup Administration Console**.
- 2 Click **NetBackup Management > Policies** in the left pane.
- 3 In the **All Policies** pane, right-click and create a new one.
- 4 From the **Policy type** list, select **NAS-Data-Protection**
- 5 From the **Policy Storage** list, select **Storage Lifecycle Policy** for backup from snapshots.
- 6 The **Perform snapshot backups** option is selected by default.
- 7 (Optional) Select the **Take checkpoints every** check box and enter a value in minutes.

- 8 (Optional) To track changes between two snapshots, select the **Enable vendor change tracking for incremental backups** check box.

Note: At a given time, you can only use either the **Enable vendor change tracking for incremental backups** or the **Use Accelerator** check-box.

- 9 Click **Options** and set a value for the following parameters:

- **Snapshot Type**
- **Snapshot Host**


Note: The **Snapshot Type *Auto*** is supported only with NetBackup client version 10.0 onwards.

- 10 In the **Dynamic Data Streaming Attributes** section, the **Allow dynamic streaming** option is selected by default. Configure the following attributes:

Note: With this setting, the **Allow multiple data streams** option is also selected.

- **Maximum number of streams per volume**
The number of streams per volume must be between 1 to 20. The default value is 4.
- **Maximum number of files in a batch**
The number of files in a batch must be between 1 to 2000. The default value is 300.

11 (Optional) Select the **Use Accelerator** check box.



Note: At a given time, you can use either **Enable vendor change tracking for incremental backups** or the **Use Accelerator** check box.

12 On the **Schedules** tab, set the required attributes for the schedule.

Note: For Instant Recovery, only Snapshots and copy snapshots to a storage unit is supported and is selected by default.


13 On the **Clients** tab, from the **NAS Vendor** list, select the preferred vendor.

14 Click **New** to add a new client.

15 On the **Backup Selections** tab, select the preferred protocol.


- NFS
- SMB

- 16** (Optional) If you want to want to use volumes that support both NFS and SMB protocol, select the **Include Mixed Volume** check box.



- 17** From the **Backup Host Pool** list, select the preferred pool and click **New**. If no backup host pools of the relevant type are available, you can see a dialog prompting you to create one, click **Yes** in the dialog to create one.


Note: If Backup host pool contains any backup host older than NetBackup 10.0 and **Auto** snapshot type is selected, the job may fail.



- 18** In the **Add Backup Selection** dialog box, click **Browse**.

If you add a subdirectory from volume in the backup selection then the policy validation fails.

- 19 In the **NAS Assets Selection** section, select the preferred volumes and click **OK**.



- 20 On the **Exclude Volumes** tab, in the **Volume to exclude** field, add the preferred volumes that you don't want to backup.
- 21 Click **OK**.

About restores from NAS-Data-Protection policy

Using the Backup, Archive, and Restore interface, you can restore individual directories or files (or an entire NAS volume) from a DNAS backup. The procedure is the same as that for restoring from a regular backup as described in the *NetBackup Backup, Archive, and Restore Getting Started Guide*.

Types of restore:

- **Point in Time Rollback (PIT):** Supported only for the NAS array vendors that provides PIT rollback capability.
- **Restore using snapshot image:** Restore data at alternate location on any supported file system and a connected NetBackup host.
- **Restore using backup image:** You can restore data at alternate location on any supported file system and a connected NetBackup host.

Note the following important points for restore:

- Original location restore is not supported for NAS-Data-Protection policy. We only support restore to alternate location. As the policy does not require the volumes to be mounted on any client as a pre-requisite, there might not be any original location. If you want to restore to the original volume, you can mount the original volume on a NetBackup host and then specify that volume at the time of restore.
- The destination client for restore must be a NetBackup host.
- If you select either of the following rename options, ensure that you change the destination path:
 - **Rename hard links**
 - **Rename soft links**

FlashBackup configuration

This chapter includes the following topics:

- [About FlashBackup](#)
- [FlashBackup restrictions](#)
- [Configuring a FlashBackup policy](#)
- [Configuring FlashBackup policy for backward compatibility \(UNIX only\)](#)

About FlashBackup

FlashBackup is a policy type that combines the speed of raw-partition backups with the ability to restore individual files. The features that distinguish FlashBackup from other raw-partition backups and standard file system backups are these:

- Increases backup performance as compared to standard file-ordered backup methods. For example, a FlashBackup of a file system completes faster than other types of backup in the following case:
 - the file system contains a large number of files
 - most of the file system blocks are allocated
- Enables restore of individual files from raw-partition backups.
- Backs up the following file systems: VxFS (Solaris, HP, Linux, AIX), ufs (Solaris), Online JFS (HP), and NTFS (Windows).
See [“FlashBackup restrictions”](#) on page 72.
- Supports multiple data streams, to further increase the performance of raw-partition backups when multiple devices are in the **Backup Selections** list.
- For a complete list of supported platforms, snapshot methods, and data types, see the NetBackup Snapshot Client Compatibility document:
<http://www.netbackup.com/compatibility>

FlashBackup restrictions

Note the following restrictions:

- FlashBackup policies do not support file systems that HSM manages.
- FlashBackup policies for UNIX clients do not support Instant Recovery.
- FlashBackup does not support VxFS storage checkpoints that the VxFS_Checkpoint snapshot method uses.
- FlashBackup supports the following I/O system components: ufs, VxFS, and Windows NTFS file systems, VxVM volumes and LVM volumes, and raw disks. Other components (such as non-Veritas storage replicators or other non-Veritas volume managers) are not supported.
- FlashBackup on Linux supports only the VxFS file system on VxVM volumes. For Linux clients, no other file system is supported, and VxFS file systems are not supported without VxVM volumes.
- FlashBackup on AIX supports only the VxFS file system, with VxVM or LVM volumes. For AIX clients, no other file system is supported, and the data must be over a VxVM or LVM volume.
- Note these restrictions for Windows clients:
 - The use of FlashBackup in a Windows Server Failover Clustering (WSFC) environment is supported, with the following limitation: Raw partition restores can only be performed when the disk being restored is placed in extended maintenance mode or removed from the WSFC resource group.
 - FlashBackup-Windows and Linux policies do not support a Client Direct restore.
 - FlashBackup-Windows policies support Instant Recovery, but only for backup to a storage unit (not for snapshot-only backups).
 - FlashBackup-Windows policies do not support the backup of Windows system-protected files (the System State, such as the Registry and Active Directory).
 - FlashBackup-Windows policies do not support the backup of Windows OS partitions that contain the Windows system files (usually C:).
 - FlashBackup-Windows policies do not support the backup of Windows System database files (such as RSM Database and Terminal Services Database).
 - FlashBackup-Windows policies do not support "include" lists (exceptions to client "exclude" lists).

- A restore from a FlashBackup-Windows image will not work if the destination client is not a Windows host.
- FlashBackup-Windows supports the backup and restore of NTFS files that are compressed. FlashBackup-Windows does support Windows NTFS encryption and compression, but not NetBackup's compression or encryption.

Note: The compressed NTFS files are backed up and restored as compressed files (they are not uncompressed).

- FlashBackup (UNIX) does not support any type of compression or encryption, whether the encryption/compression is set in the NetBackup policy or in the Operating System.

Restores of Windows encrypted files and hard links

When restoring individual files from FlashBackup-Windows images, note: if the files being restored are encrypted or are hard links, the NetBackup Client Service must be logged on as Administrator. Under services on the control panel, change the log on for the NetBackup Client Services from **Local System Account** to **Administrator**.

Configuring a FlashBackup policy

Use the following procedure.

To configure a FlashBackup policy

- 1 On the master server, open the NetBackup Administration Console.
- 2 Click **Policies**.

Change Policy - Flash_Backup_policy1

Server: nbu-id1

Attributes Schedules Clients Backup Selections

Policy type: FlashBackup-Windows

Destination:

Data classification: <No data classification>

Policy storage: Any_available

Policy volume pool: NetBackup

☐ Take checkpoints every: 0 minutes

☐ Limit jobs per policy:

Job priority: 0 (higher number is greater priority)

Media Owner: Any

Snapshot Client and Replication Director

☐ Perform block level incremental backups

☐ Use Replication Director

☒ Perform snapshot backups Options...

☐ Retain snapshot for Instant Recovery or SLP management

☐ Hyper-V server:

☒ Perform off-host backup

Use: Alternate client

Machine:

☒ Go into effect at: 02/23/2015 18:46:18

☐ Follow NFS

☐ Cross mount points

☐ Compress

☐ Encrypt

Collect disaster recovery information for:

☐ Bare Metal Restore

☐ Collect true image restore information

☐ With move detection (Required for synthetic backups and Bare Metal Restore)

☐ Allow multiple data streams

☐ Disable client-side deduplication

☐ Enable granular recovery

☐ Use Accelerator

☐ Enable optimized backup of Windows deduplicated volumes

Keyword phrase (optional):

Microsoft Exchange Server Attributes

Exchange DAG or Exchange 2007 replication (LCR/CCR)

Database backup source:

Preferred server list... (Exchange DAG only)

OK Cancel Help

- 3 In the **All Policies** pane, right-click and select **New Policy...** to create a new policy.

- 4 On the **Attributes** tab, select the Policy type: **FlashBackup** for UNIX clients, or **FlashBackup-Windows** for Windows clients.

FlashBackup-Windows supports the backup and restore of NTFS files that are compressed.

The files are backed up and restored as compressed files (they are not uncompressed).

- 5 Specify the storage unit.

FlashBackup and FlashBackup-Windows policies support both tape storage units and disk storage units.

- 6 Select a snapshot method in one of the following ways:

- Click **Perform snapshot backups** on the **Attributes** tab.
For a new policy, NetBackup selects a snapshot method when the backup starts.
For a copy of a policy that was configured for a snapshot method, click the **Snapshot Client Options** option and set the method to **auto**. NetBackup selects a snapshot method when the backup starts.
- Click **Perform snapshot backups**, click the **Snapshot Client Options** option and select a snapshot method.
See [“Selecting the snapshot method”](#) on page 43.

- 7 Windows only: to enable the backup for Instant Recovery, select **Retain snapshots for Instant Recovery or SLP management**.

Instant Recovery is not supported for FlashBackup with UNIX clients.

- 8 UNIX only: if you selected nbu_snap or VxFS_Snapshot as the snapshot method, specify a raw partition as cache, in either of these ways:

- Use the Host Properties node of the Administration Console to specify the default cache device path for snapshots. Click **Host Properties > Clients**, select the client, then **Actions > Properties, UNIX Client > Client Settings**.
- Use the **Snapshot Client Options** dialog to specify the cache.
See [“Entering the cache”](#) on page 97.
The partition to be used for the cache must exist on all clients that are included in the policy.

- 9 To shift backup I/O to an alternate client, or to a NetBackup media server or third-party copy device (UNIX clients only), select **Perform off-host backup**.

For FlashBackup, the **Use data mover** option is supported for UNIX clients only.

- 10** To reduce backup time when more than one raw partition is specified in the **Backup Selections** list, select **Allow multiple data streams**.
- 11** Use the **Schedules** tab to create a schedule.

FlashBackup policies support full and incremental types only. User backup and archive schedule types are not supported.

A full FlashBackup backs up the entire disk or raw partition that was selected in the **Backup Selections** tab (see next step). An incremental backup backs up individual files that have changed since their last full backup, and also backs up their parent directories. The incremental backup does not back up files in parent directories unless the files have changed since the last full backup.

For incremental backups, a file is considered “changed” if its **Modified Time** or **Create Time** value was changed.

Note on FlashBackup-Windows: The NTFS Master File Table does not update the **Create Time** or **Modified Time** of a file or folder when the following changes are made:

- Changes to file name or directory name.
- Changes to file or directory security.
- Changes to file or directory attributes (read only, hidden, system, archive bit).

- 12** On the **Backup Selections** tab, specify the drive letter or mounted volume (Windows) or the raw disk partition (UNIX) containing the files to back up.

For Windows

```
Windows examples    \\.\\E:
                   \\.\\E:\mounted volume\
```

where:

- `\\.\E:` is a Windows disk volume mounted on a drive letter.
- `\\.\E:\mounted_volume\` (note the trailing backslash) is a Windows disk volume without a drive letter mounted on a directory (Windows repare point). The drive must be designated exactly as shown (`E:\` is not correct). Backing up the drive containing the Windows system files (usually the C drive) is not supported.

```
Solaris examples      /dev/rdisk/c1t0d0s6
                     /dev/vx/rdisk/volgrp1/vol1
```

HP examples /dev/rdisk/clt0d0
 /dev/vx/rdsk/volgrpl/voll

For UNIX

The

Backup Selections tab must specify the raw (character) device corresponding to the block device over which the file system is mounted. For example, to back up /usr, mounted on /dev/dsk/clt0d0s6, enter raw device /dev/rdsk/clt0d0s6. **Note the r** in /rdsk.

Note: Wildcards (such as /dev/rdsk/c0*) are not permitted. Specifying the actual device file name such as /devices/pci@1f,0/pci@1/scsi@3/sd@1,0:d,raw is not supported.

Note: Wildcards (such as /dev/rdsk/c0*) are not permitted. Specifying the actual device file name such as /devices/pci@1f,0/pci@1/scsi@3/sd@1,0:d,raw is not supported.

13 Use the **Clients** tab to select clients to back up.

Each client in the client list must contain all the raw partitions that are specified in the **Backup Selections** tab. Exclude and include lists cannot be used to circumvent this requirement.

14 On the policy **Attributes** tab: if you click **Apply** or **OK**, a validation process checks the policy and reports any errors. If you click **Close**, no validation is performed.

Configuring FlashBackup policy for backward compatibility (UNIX only)

Before NetBackup 5.0, FlashBackup was a separate product with two built-in snapshot methods: snapctl driver (nbu_snap) for Solaris clients and a VxFS snap driver for HP clients. The configuration procedure for a pre-5.0 FlashBackup policy was different from the procedure in 5.0 and later, as follows:

- Unless FlashBackup was combined with NetBackup 4.5 ServerFree Agent, the snapshot methods were preset to the snapctl driver (Solaris) and VxFS snap driver (HP).
- The cache partition for the snapctl driver and VxFS snap driver had to be specified as a CACHE= entry in the policy's file list.

- To use multiple data streams, other directives had to be added to the policy's **Backup Selections** (file) list.

The following procedure and related topics explain how to configure a FlashBackup policy with a `CACHE=` entry in the policy's **Backup Selections** list. This means of configuration is provided for backward compatibility.

To configure FlashBackup policy for backward compatibility (UNIX only)

- 1 Leave **Perform snapshot backups** deselected on the policy **Attributes** tab. NetBackup uses `nbu_snap` (snapctl driver) for Solaris clients or `VxFS_Snapshot` for HP.
- 2 On the policy's **Backup Selections** tab, specify at least one cache device by means of the `CACHE` directive. For example:

```
CACHE=/dev/rdisk/c0t0d0s1
```

This cache partition is for storing any blocks that change in the source data while the backup is in progress. `CACHE=` must precede the source data entry. Note the following:

- Specify the raw device, such as `/dev/rdisk/c1t0d0s6`. Do not specify the block device, such as `/dev/dsk/c1t0d0s6`.
- Do not specify the actual device file name. For example, the following is not allowed:

```
/devices/pci@1f,0/pci@1/scsi@3/sd@1,0:d,raw
```

- Wildcards (such as `/dev/rdisk/c0*`) are not allowed.
- The `CACHE` entry must precede the entry for the source data you want to back up.
- All entries in the **Backup Selections** list, including the source data, must be the full path name of a raw device in the form:

On Solaris: `/dev/rdisk/cxtxdxsx`

On HP: `/dev/rdisk/cxtxdx`

where x is an integer.

- For multiple data streams, you can include multiple entries in the **Backup Selections** list.

For example:

```
CACHE=/dev/rdisk/c1t4d0s0
```

```
/dev/rdisk/c1t4d0s7
```

```
CACHE=/dev/rdisk/c1t4d0s1
```

```
/dev/rdisk/c1t4d0s3  
/dev/rdisk/c1t4d0s4
```


- See “Requirements for the cache partition” on page 80.
- See “About the cache partition” on page 79.

About the cache partition

The snapctl driver and VxFS snap driver are copy-on-write snapshot methods that require a cache partition. For FlashBackup clients before NetBackup 5.0, the cache partition was specified in the policy's **Files** tab as a **CACHE = raw_partition** entry. (The **Files** tab is now called the **Backup Selections** tab.)

Note: CACHE entries are allowed only when the policy's **Perform snapshot backups** option is deselected. If **Perform snapshot backups** is selected, NetBackup attempts to back up the CACHE entry and the backup fails.

Figure 5-1 Backup Selections list with CACHE entry



All entries must specify the raw device, such `/dev/rdisk/c0t0d0s1`. Do not use the actual file name; you must use the link form of `ctxdxsx`.

Requirements for the cache partition

Note the following requirements:

- Must reside on the same host as the raw partitions containing the source data to back up.
- Cannot be the raw partition being backed up.
- Cannot be a mounted file system. Any data that is configured on this device may be overwritten by the copy-on-write process.
- On Solaris, the same cache partition may be used simultaneously by multiple backups (two policies can use the same cache partition at the same time). On HP, multiple backups cannot use the same cache partition simultaneously. If multiple policies list the same cache partition on HP systems, backups naming those policies must run at different times to prevent a failure.
- The cache partition must have enough space to hold all the writes to the source data that may occur during the backup. Backups during off-peak hours normally require a smaller cache than those during peak activity.
See [“Determining a size for the cache partition”](#) on page 95.

Directives for multiple data streams

For multiple data streams, certain directives must be added to the policy’s **Backup Selections** tab.

- The number of backups that are started depends on the directives in the **Backup Selections** tab.
- The maximum number of concurrent backups depends on the number of available drives in the storage units and the maximum jobs parameters. An example of a maximum jobs parameter is **Limit jobs per policy**.

Note: Only one data stream is created for each physical device on the client. You cannot include the same partition more than once in the **Backup Selections** list.

The directives that you can use in the **Backup Selections** list for a FlashBackup policy are as follows:

- `NEW_STREAM`
- `CACHE=value` (the `CACHE` directive is required)
See [“Requirements for the cache partition”](#) on page 80.
- `UNSET`
- `UNSET_ALL`

Each backup begins as a single stream of data. The start of the **Backup Selections** list up to the first `NEW_STREAM` directive (if any) is the first stream. Each `NEW_STREAM` entry causes NetBackup to create an additional stream or backup.

Note that all file paths that are listed between `NEW_STREAM` directives are in the same stream.

Table 5-1 shows a **Backup Selections** list that generates four backups:

Table 5-1 Example of Backup Selections list

Example	On Solaris systems:	On HP systems:
1	CACHE=/dev/rdisk/clt3d0s3 /dev/rdisk/clt0d0s6	CACHE=/dev/cache_group/rvol1c /dev/vol_grp/rvol1
2	NEW_STREAM /dev/rdisk/clt1d0s1	NEW_STREAM UNSET CACHE CACHE=/dev/cache_group/rvol2c /dev/vol_grp/rvol2
3	NEW_STREAM UNSET CACHE CACHE=/dev/rdisk/clt3d0s4 /dev/rdisk/clt2d0s5 /dev/rdisk/clt5d0s0	NEW_STREAM UNSET CACHE CACHE=/dev/cache_group/rvol3c /dev/vol_grp/rvol3 /dev/vol_grp/rvol3a
4	NEW_STREAM UNSET CACHE CACHE=/dev/rdisk/c0t2d0s3 /dev/rdisk/clt6d0s1	NEW_STREAM UNSET CACHE CACHE=/dev/cache_group/rvol4c /dev/vol_grp/rvol4

The backup streams are issued as follows. The following items correspond in order to the numbered items in Table 5-1:

1. The first stream is generated automatically and a backup is started for /dev/rdisk/clt0d0s6 (Solaris) or /dev/vol_grp/rvol1 (HP). The `CACHE=` entry sets the cache partition to /dev/rdisk/clt3d0s3 (Solaris) or /dev/cache_group/rvol1c (HP).
2. The first `NEW_STREAM` directive (2) starts a second stream to back up /dev/rdisk/clt1d0s1 (Solaris) or /dev/vol_grp/rvol2 (HP). On Solaris systems, this backup uses the same cache partition. On HP systems, a different cache partition must be defined for each stream (`CACHE=/dev/cache_group/rvol2c`).

3. The second `NEW_STREAM` directive (3) starts a backup for `/dev/rdisk/clt2d0s5` and `/dev/rdisk/clt5d0s0` (Solaris) or `/dev/vol_grp/rvol3` and `/dev/vol_grp/rvol3a` (HP). These two partitions are backed up serially within the stream. In addition, the `UNSET CACHE` directive unsets the previous cache setting and the `CACHE=` directive sets a new cache partition for this backup.
4. The last `NEW_STREAM` directive (4) starts a backup for `/dev/rdisk/clt6d0s1` (Solaris) or `/dev/vol_grp/rvol4` (HP). Like the third stream, this one also unsets the cache directive and defines a new cache partition.

Policy-specific directives such as `CACHE` are passed to the client with the current stream and all subsequent streams, until the directive is redefined or unset.

If the directive is encountered again, its value is redefined.

An `UNSET` or `UNSET_ALL` directive unsets a directive that was previously defined in the **Backup Selections** list. Note the following:

- `UNSET` unsets a policy-specific directive so it is not passed with additional streams. The directive that was unset can be defined again later in the **Backup Selections** list to be included in the current or a later stream.
- `UNSET_ALL` has the same effect as `UNSET` but affects all policy-specific directives that have been defined up to this point in the **Backup Selections** list. If you use it, `UNSET_ALL` must appear immediately after the second or later `NEW_STREAM` directive.

Instant Recovery configuration

This chapter includes the following topics:

- [About Instant Recovery capabilities](#)
- [Instant Recovery requirements](#)
- [Instant Recovery restrictions](#)
- [Giving full server privileges to the media server](#)
- [About Instant Recovery](#)
- [Configuring a policy for Instant Recovery](#)
- [About sizing the cache for Instant Recovery copy-on-write snapshots](#)
- [Instant Recovery for databases](#)
- [About storage lifecycle policies for snapshots](#)

About Instant Recovery capabilities

The Instant Recovery feature of Snapshot Client enables high-speed data retrieval from disk by means of the standard NetBackup user interface. Note the following features:

- Uses snapshot technologies to create disk images.
- Can create a snapshot and a backup to tape or disk, from one policy.
- Enables random-access (non-sequential) restores of dispersed files from full backups.

- Can restore to a different path or host.
- Provides the resource management by means of a rotation schedule.

Instant Recovery requirements

Note the following requirements:

Instant Recovery restrictions

Note the following restrictions:

- Instant Recovery snapshots must not be manually removed or renamed, otherwise the data cannot be restored.
- Any media server that is used in an Instant Recovery backup must have full server privileges.


See [“Giving full server privileges to the media server”](#) on page 84.

Giving full server privileges to the media server

A media server used in an Instant Recovery backup must have full server privileges. If it does not have server privileges, the snapshots created by each backup do not properly expire in the NetBackup catalog.

To give full server privileges to the media server

- 1 In the NetBackup Administration Console, click **Host Properties > Master Server > *double click master server* > Master Server Properties > Servers**.



- 2 Make sure that the media server is listed under **Additional Servers**, not under **Media Servers**.


Note: on UNIX, this procedure places a `SERVER = host` entry in the `bp.conf` file for each host that is listed under **Additional Servers**. In the `bp.conf` file, the media server must not be designated by a `MEDIA_SERVER = host` entry.

About Instant Recovery

Standard NetBackup can use disks for backup and restore. The Instant Recovery feature of Snapshot Client extends that capability by exploiting the speed of snapshots. A snapshot is created with minimal effect on client access to data or on the speed of client transactions. If snapshots are made frequently, you can restore an accidentally deleted file in a matter of seconds.

Figure 6-1 shows an example.

Figure 6-1 Instant Recovery from snapshot on disk



The on-disk snapshots become point-in-time versions of file systems or volumes, to be kept or discarded as needed. The data can be restored from local disk; no need to mount a tape or other remote storage.

NetBackup Instant Recovery creates snapshot A of the client data on disk. One hour later, as scheduled, NetBackup creates snapshot B, also on disk, followed one hour later with snapshot C. By using the appropriate snapshot, you can restore data directly from disk.

Note: NetBackup Instant Recovery retains the snapshot. The snapshot can be used for restore even if the client has been restarted.

The following sections provide some background information.

About snapshot and backup for Instant Recovery

About NetBackup catalog maintenance

For Instant Recovery backups, NetBackup automatically updates the catalog to keep it correlated with the snapshots on the client. If not kept up-to-date, the catalog might refer to snapshots that no longer exist on the client, due to user activity (snapshot replacement or deletion).

NetBackup includes a maintenance program (`bpfficorr`) that runs after backups and restores. It can also be run manually to update the catalog if an Instant Recovery snapshot is accidentally deleted or renamed.

For man page information on the `bpfficorr` command, refer to the [NetBackup Commands guide](#).

About snapshot management

Because snapshots require disk space, they cannot be retained forever. To balance the space consumed against the convenience of having multiple snapshots available for instant recovery, you can specify how many snapshots to retain.

Means of controlling snapshots

Maximum Snapshots parameter

The **Maximum Snapshots** parameter on the **Snapshot Options** dialog box sets the maximum number of Instant Recovery snapshots to be retained at one time.

When the maximum is reached, the next snapshot causes the oldest job-complete snapshot to be deleted.

A snapshot job is considered to be complete once all its configured dependent copies (for example, Backup from Snapshot, Index, Replication) are complete.

For Instant Recovery backups, it is good practice to set the backup retention level to infinite. A retention period that is too short can interfere with maintaining a maximum number of snapshots for restore.

Configuring a policy for Instant Recovery

This section explains how to configure a policy for the backups that are enabled for Instant Recovery.

To configure a policy for Instant Recovery

- 1 On the master server, open the NetBackup Administration Console.
- 2 Click **Policies**. In the **All Policies** pane, open a policy or create a new one.
- 3
- 4 Select a storage unit (disk or tape).

If you select **Snapshots only** on the **Schedules** tab, the storage unit is not used. NetBackup creates a snapshot only.
- 5 Select **Perform snapshot backups**.
- 6 Select **Retain snapshots for instant recovery or SLP management**.

NetBackup retains the snapshot so that Instant Recovery can be performed from the snapshot. A normal backup to storage is also performed, if you do not select **Snapshots only** on the **Schedules** tab.
- 7 To save the settings, click **Apply**.

8

9 To configure a schedule, use the **Schedule** tab.

- For a snapshot only, select **Snapshots only**.
If **Snapshots only** is selected, the snapshot is not backed up to tape or other storage. NetBackup creates a snapshot on disk only. Note that you must deselect **Snapshots only** if you want to deselect **Retain snapshots for instant recovery and SLP management** on the policy **Attribute** tab. If **Snapshots and copy snapshots to a storage unit** is selected, NetBackup creates (and retains) a snapshot and backs up the client's data to the storage unit that is specified in the policy.
- You can select the retention period for snapshots under **Retention**.
- Make other schedule selections as desired, and click **OK**.

10 To enter the files and folders to be backed up, use the **Backup Selections** tab.

- When backing up Oracle database clients, refer to the [NetBackup for Oracle System Administrator's Guide](#) for instructions.

11 To specify clients to be backed up by this policy, use the **Clients** tab.

12 On the policy **Attributes** tab: if you click **Apply** or **OK**, a validation process checks the policy and reports any errors. If you click **Close**, no validation is performed.

About sizing the cache for Instant Recovery copy-on-write snapshots

A copy-on-write snapshot requires cache space for storing the changes that occur on the source device during the life of the snapshot. While the snapshot is active, any blocks that are about to be changed by user activity are copied to the cache. Blocks that do not change on the source are not copied. Compared to a fully-allocated snapshot (clone or mirror), a copy-on-write snapshot may consume relatively little disk space and can be executed very quickly.

As a rule, the appropriate size for the cache depends on the amount of user activity that occurs during the life of the snapshot. The more changes in the source data, or the longer the life of the snapshot, the more blocks that are likely to be changed. As a result, more data must be stored in the cache.

The size of the file system or raw partition does not determine cache size. If little change activity occurs on the source during the life of the snapshot, little cache space is required, even for a large file system.

Note: If the cache runs out of space, the snapshot may fail.

Cache size during restore

The size of the cache may need to increase when you restore large amounts of data from an Instant Recovery snapshot. For instance, if you restore an entire raw partition to a source device that has an active Instant Recovery snapshot, writing the restore data causes the active snapshot to cache all blocks from the raw partition that resides on the source device. A cache that was large enough when the snapshot was first initiated may no longer be large enough. The large restore multiplies the write-to-cache activity, increasing the space that is needed for cache.

Setting an adequate size for the snapshot cache

The cache must be sized appropriately in either of the following cases:

- If a lot of write activity is expected in the source data during the life of the copy-on-write snapshot.
- If you expect to restore large amounts of data from an Instant Recovery snapshot.

To set an adequate size for the snapshot cache

Use the following formula:

For raw partitions: $\text{Cache size} = \text{volume size} * \text{the number of retained snapshots}$

For file systems: $\text{Cache size} = (\text{consumed space} * \text{the number of retained snapshots})$
+ approximately 2% to 5% of the consumed space in the file system

Note:

- Consumed space means the total size of the data allocated to files in the file system, not the total size of the mounted volume.
- The additional space (2% to 5%) represents a safety measure and may not be required.

Instant Recovery for databases

To configure an Instant Recovery policy for database clients, refer to the appropriate NetBackup database agent guide.

About storage lifecycle policies for snapshots

A storage lifecycle policy is a storage plan for a set of backups. NetBackup uses the lifecycle policy to determine where to store additional copies of the backup images and how long to retain those copies. In general, short-term copies can be kept on disk (for quick restore) and long-term copies can be kept on tape or other storage.

NetBackup can manage snapshot-based backups for Instant Recovery through storage lifecycle policies.. The Instant Recovery feature makes snapshots available for quick data recovery from disk. Lifecycle policies support a lifecycle storage plan for the storage unit copies that are made during an Instant Recovery backup.

Configuring a storage lifecycle policy to manage snapshot-based backups for Instant Recovery

This section describes how to create a storage lifecycle policy to manage snapshot-based backups for Instant Recovery. The procedure focuses on snapshot-related details only.

Full procedures are available for creating storage lifecycle policies:

See the [NetBackup Administrator's Guide, Volume I](#).

To configure a storage lifecycle policy to manage snapshot-based backups for Instant Recovery

- 1 Create a lifecycle policy with two or more storage destinations.
Use the **Storage > Storage Lifecycle Policies** node of the NetBackup Administration Console. Click **Actions > New > Storage Lifecycle Policies**. Click **Add**.
 - For snapshots, select **Snapshot** on the **New Storage Destination** dialog. You can specify a retention period appropriate for snapshots (such as two weeks). Click **OK**.
 - For backup copies to disk, select **Backup** on the **New Storage Destination** dialog. Specify a disk storage unit and a longer retention period (such as six months). Click **OK**.
 - For backup copies to tape, select **Duplication** on the **New Storage Destination** dialog. Specify a tape storage unit and a longer retention period (such as five years). Click **OK** and finish creating the lifecycle policy.
- 2 Create a policy for snapshots. (Use the **Policies** node of the Administration Console.)
On the policy **Attributes** tab:

- You can specify the lifecycle policy in the **Policy storage unit / lifecycle policy** field. You can later change the lifecycle policy in the schedule, as explained later in this procedure.
- Select **Perform snapshot backups**.
- On the **Snapshot Options** dialog box, the **Maximum Snapshots (Instant Recovery only)** parameter sets the maximum number of snapshots to be retained at one time. When the maximum is reached, the next snapshot causes the oldest job-complete snapshot to be deleted.

A snapshot is considered to be job complete once all its configured dependent copies (for example, Backup from Snapshot, Index, Replication) are complete.

Note that if you also set a snapshot retention period of less than infinity in the lifecycle policy, the snapshot is expired when either of these settings takes effect (whichever happens first). For example, if the **Maximum Snapshots** value is exceeded before the snapshot retention period that is specified in the lifecycle policy, the snapshot is deleted.

Policy validation fails if there is a mismatch of retention found on the snapshot. For example, if the **Maximum Snapshots (Instant Recovery only)** parameter is set to any value other than Managed by SLP and the SLP used in the same policy has **Fixed** retention for the Snapshot job the policy validation fails. If you have such a policy configured on a pre- 7.6 NetBackup master server, it is advisable that you validate and correct the policy after you upgrade to a NetBackup 10.0 master server.

3 Create a schedule for the policy.

You can create a single schedule for backups and let the lifecycle policy govern their destinations and retention periods, as follows:

- Under **Destination**, if you selected **Retain snapshots for Instant Recovery and SLP management** on the policy **Attributes** tab, make sure that **Snapshots and copy snapshots to a storage** unit is selected on the schedule (not **Snapshots only**).
Important: if you select **Snapshots only** on the schedule, a lifecycle policy cannot be used.
- In the **Override policy storage selection** field, select the lifecycle policy that you created in [1](#).
- Under **Schedule type**, set an appropriate frequency, such as 1 day.

When the Snapshot Client policy executes this schedule, the lifecycle policy named in the **Override policy storage selection** field creates images on the destinations that are named in the lifecycle policy. The lifecycle policy also sets

the retention periods for the images it creates. In this example, the retention is six months for backups to disk and five years for tape.

Storage lifecycle policies and Snapshot Client troubleshooting

The section includes information about various error messages related to SLPs and snapshots.

If you configure a snapshot method for a policy, and the schedule specifies a lifecycle policy, the life cycle policy must include a snapshot destination. Otherwise, an error such as the following appears in the NetBackup Problems report:

```
snapshot backup: tashina11_1204305543 cannot be used with a  
lifecycle policy NoSnapshot that does not include a snapshot  
destination.
```

Configuration of software-based snapshot methods

This chapter includes the following topics:

- [Software-based snapshot methods](#)

Software-based snapshot methods

This topic provides notes and configuration assistance for the Snapshot Client methods exclusive of those designed for disk arrays.

For configuration notes on disk array methods, see the following:

About nbu_snap

The **nbu_snap** snapshot method is for Solaris clients only. It is for making copy-on-write snapshots for UFS or VxFS file systems.

The information in this section applies to either Standard or FlashBackup policy types.

nbu_snap is not supported in clustered file systems. It is not supported as the selected snapshot method or as the default `snapctl` driver when you configure FlashBackup in the earlier manner.

See [“Configuring FlashBackup policy for backward compatibility \(UNIX only\)”](#) on page 77.

An alternative copy-on-write snapshot method for clustered file systems is VxFS_Snapshot with a FlashBackup policy.

nbu_snap does not support VxVM volumes that belong to a shared disk group.

Cache device requirements

Note the following:

- The cache device is a raw disk partition: either a logical volume or physical disk. The cache is used for storing the portions of the client's data that incoming write requests change while the copy-on-write is in progress.
- For the cache device, do not select an active partition containing valuable data. Any data in that partition is lost when the snapshot is complete.

Warning: Choose a cache partition carefully! The cache partition's contents are overwritten by the snapshot process.

- Specify the raw partition as the full path name of either the character special device file or the block device file. For example:

Solaris raw partition: `/dev/rdisk/c2t0d3s3`

Or

`/dev/dsk/c2t0d3s3`

VxVM volume: `/dev/vx/rdsk/diskgroup_1/volume_3`

Or

`/dev/vx/dsk/diskgroup_1/volume_3`

Note: Do not specify wildcards (such as `/dev/rdsk/c2*`) as paths.

- The cache partition must be unmounted.
- The cache partition must reside on the same host as the snapshot source (the client's data to back up).
- The partition must have enough space to hold all the writes to the partition that may occur during the backup. Note that backups during off-peak periods normally require a smaller cache than a backup during peak activity.
See [“Determining a size for the cache partition”](#) on page 95.

- For the Media Server or Third-Party Copy Device method: the host containing the snapshot source and cache must be visible to the media server or third-party copy device.
- For the Media Server or Third-Party Copy Device method: the disk containing the cache must meet certain requirements.
See [“Disk requirements for Media Server and Third-Party Copy methods”](#) on page 152.

Determining a size for the cache partition

The required size for the cache partition depends on user write activity during the backup, not on the size of the client's file system. If the backup occurs when user activity is heavy, a larger cache is required.

To determine a size for the cache partition

- 1 Consider the period in which the backup is scheduled to occur: the more user activity that is expected, the larger the cache required.

You should execute the following procedure at an appropriate period, when your snapshot backups typically run. If user activity at your site is known to vary with the time of day, a different time could bring very different results.

- 2 Make sure that a raw partition is available on a separate disk.
See [“Cache device requirements”](#) on page 94.
- 3 During the appropriate backup period, create an `nbu_snap` snapshot by entering the following as root:

```
/usr/opensv/netbackup/bin/driver/snapon snapshot_source cache
```

where *snapshot_source* is the partition on which the client's file system is mounted, and *cache* is the raw partition to be used as copy-on-write cache. For example:

```
/usr/opensv/netbackup/bin/driver/snapon /omo_cat3  
/dev/vx/rdisk/zeb/cache
```

Example output:

```
matched /omo_cat3 to mnttab entry /omo_cat3  
mount device: /dev/vx/dsk/omo/vol03 fstype: vxfs  
snapshot 29 enabled on /omo_cat3 at 06/05/03 15:16:02
```

- 4 In `/usr/opensv/netbackup/bin/driver`, enter the `snaplist` and `snappchelist` commands.

`snaplist` shows the following:

- Id of each snapshot
- Size of the partition containing the client file system
- Amount of file system write activity in 512-byte blocks that occurred during the `nbu_snap` snapshot (under the `cached` column).
The more blocks that are cached as a result of user activity, the larger the cache partition required.

`snapcachelist` shows each cache device in use and what percentage has been used (`busy`). For each cache device that is listed, `busy` shows the total space that is used in the cache. This value indicates the size of the raw partition that may be required for `nbu_snap` cache.

More details are available on the `snap` commands.

See [“nbu_snap commands”](#) on page 207.

The `snap` commands can be used in a script.

If the cache partition is not large enough, the backup fails with status code 13, “file read failed.” The `/var/adm/messages` log may contain errors such as the following:

```
Mar 24 01:35:58 bison unix: WARNING: sn_alloccache: cache
/dev/rdsk/c0t2d0s3 full - all snaps using this cache are now
unusable
```

5 Using the information that `snaplist` and `snapcachelist` provide, you have several options:

- Specify a larger (or smaller) partition as cache, depending on the results from `snaplist` and `snapcachelist`.
- Reschedule backups to a period when less user activity is expected.

- If multiple backups use the same cache, reduce the number of concurrent backups by rescheduling some of them.
- 6 When you are finished with the snapshot, you can remove it by entering the following:

```
/usr/opensv/netbackup/bin/driver/snapoff snapid
```

where *snapid* is the numeric id of the snapshot that was created earlier.

NetBackup policies do not control any snapshots that you create manually with the `snapon` command. When `snapon` is run manually, it creates a copy-on-write snapshot only. The snapshot remains on the client until it is removed by entering `snapoff` or the client is restarted.

Entering the cache

For the `nbu_snap` and `VxFS_Snapshot` methods, you must identify a raw partition that the copy-on-write process uses, in any of the following ways.

To enter the cache

- 1 If manually selecting the snapshot method on the **Snapshot Client Options** dialog, you have two options for specifying the raw cache partition:
 - Under **Host Properties > Clients > Client Properties dialog > UNIX Client > Client Settings**, specify the raw partition in the **Default cache device path for snapshots** field. For example: `/dev/rdisk/c1t0d0s6`. This setting applies to the client in all policies.
 - Or, under **Policies > Attributes > Snapshot Client Options** dialog, specify the raw partition in the **Cache device path** field. This cache setting applies to all clients in the current policy, and overrides the cache setting in the **Client Settings** dialog.
- 2 If you want NetBackup to select the `nbu_snap` or `VxFS_Snapshot` methods by means of the auto method, specify the cache on the **Host Properties > Clients > Client Properties dialog > UNIX Client > Client Settings**.
- 3 In a FlashBackup policy: if **Perform snapshot backups** is NOT selected, you must use a `CACHE=` directive in the **Backup Selections** tab.

This cache setting applies to all clients in the current policy and overrides the cache setting in the **Host Properties** dialog. (This means of configuring the cache will be discontinued in a future release.)

About VxFS_Checkpoint

The VxFS_Checkpoint snapshot method is for making copy-on-write snapshots. This method is one of several snapshot methods that support Instant Recovery backups. Note that for VxFS_Checkpoint, the Instant Recovery snapshot is made on the same disk file system that contains the client's original data.

For VxFS_Checkpoint, VxFS 3.4 or later with the Storage Checkpoints feature must be installed on the NetBackup clients. HP requires VxFS 3.5; AIX and Linux require VxFS 4.0.

Note: On the Red Hat Linux 4 platform, VxFS_Checkpoint snapshot method supports Storage Foundation 5.0 MP3 RP3 HF9 or later versions.

Note the following:

- The VxFS_Checkpoint method is not supported for backing up raw partitions (whether **FlashBackup** or **Standard** policies).
- Make sure that enough disk space is available for the checkpoint. The file system containing the snapshot source should have at least 10% free space to successfully implement the checkpoint.

VxFS multi-volume system

VxFS_Checkpoint and VxVM are the only snapshot methods in Snapshot Client that support the multi-volume file system (MVS) feature of VxFS 4.0.

The multi-volume file system feature requires a VxVM 4.0 volume set. With volume sets, you can group related volumes into one volume set, and mount a VxFS file system on it. This means that a VxFS file system can be mounted on more than one volume. This feature allows file systems to make best use of the different performance and availability characteristics of the underlying volumes. For example, file system metadata can be stored on volumes with higher redundancy, and user data on volumes with better performance.

For background and configuration assistance with multi-volume file systems, refer to the *Veritas File System 4.0 Administrator's Guide* and the *Veritas Volume Manager 4.0 Administrator's Guide*.

Note: Off-host backup is not supported for a VxFS 4.0 multi-volume system.

Storage Checkpoint disk usage

The `ls` command does not list Storage Checkpoint disk usage. This means that the primary volume may appear to have available space even if it is full. You must use the `fscckptadm list` command to show Storage Checkpoint disk usage. Refer to the *Veritas File System Administrator's Guide* for more information on `fscckptadm`.

A new Storage Checkpoint is created whenever a VxFS_Checkpoint policy is executed.


Checkpoint retention schedules

For Instant Recovery storage checkpoints, no data is moved between the client and the media server. Instead, a VxFS Storage Checkpoint is created on the client. For Oracle clients, the file names and directory names are sent to the server for the catalog. In the case of file systems (non-database clients), only the directory name is sent, not the file data.

Changes in the files on the client's primary fileset are reflected in the Storage Checkpoint until the backup policy runs again, creating another Storage Checkpoint. Storage Checkpoints are created and retained until the maximum checkpoints threshold is exceeded, when the oldest checkpoint is removed.

Figure 7-1 shows the Instant Recovery checkpoint expiration schedule. If the maximum checkpoint value is set to 2 and a third checkpoint is created, the oldest is removed.

Figure 7-1 Instant Recovery Retention Schedule for Storage Checkpoints



A new Storage Checkpoint is created each time a backup is started.

Block-Level restore

If only a small portion of a file system or database changes on a daily basis, full restores are unnecessary. The VxFS Storage Checkpoint mechanism keeps track of the data blocks that were modified since the last checkpoint was taken. Block-level restores take advantage of this feature by restoring only changed blocks, not the entire file or database. The result is faster restores when you recover large files.

About VxFS_Snapshot

The VxFS_Snapshot method is for making copy-on-write snapshots of local Solaris or HP clients. Off-host backup is not supported with this snapshot method.

Note the following:

- VxFS_Snapshot supports the FlashBackup policy type only.
- The VxFS_Snapshot method can only be used to back up a single file system. If multiple file systems are specified in the policy's **Backup Selections** list when you use this method, the backup fails.
- In a FlashBackup policy, if the **Backup Selections** list contains CACHE= entries, FlashBackup does support the backup of multiple file systems from a single policy. For each file system, a separate cache must be designated with the CACHE= entry. Make sure you create a separate policy for each file system. See [“Configuring FlashBackup policy for backward compatibility \(UNIX only\)”](#) on page 77.
- You must designate a raw partition to be used for copy-on-write cache.
Raw partition example:

Solaris: `/dev/rdisk/clt0d0s3`

Or

`/dev/dsk/clt0d0s3`

HP: `/dev/rdisk/clt0d0`

Or

`/dev/dsk/clt0d0`

See [“Cache device requirements”](#) on page 94.

See [“Entering the cache”](#) on page 97.

- VxFS_Snapshot is the default snapshot method for FlashBackup clients running HP, when **Perform snapshot backup** is not selected for the backup policy.

About VxVM

The **VxVM** snapshot method is for making mirror snapshots with Veritas Volume Manager 3.1 or later snapshot mirrors. (On Windows, make sure that VxVM has the latest VxVM service packs and updates.)

Note: On the Red Hat Linux 4 platform, VxVM snapshot method supports Storage Foundation 5.0 MP3 RP3 HF9 or later versions.

The **VxVM** snapshot method works for any file system that is mounted on a VxVM volume. However, before the backup is performed, the data must be configured with either of the following: a VxVM 3.1 or later snapshot mirror or a VxVM 4.0 or later cache object. Otherwise, the backup fails.

Note the following:

- See [“Creating a snapshot mirror of the source”](#) on page 101.
Or refer to your *Veritas Volume Manager* documentation.
- Help is available for configuring a cache object.
See [“About VxVM instant snapshots”](#) on page 102.
Or refer to your *Veritas Volume Manager* documentation.
- For Instant Recovery backups of the data that is configured on VxVM volumes on Windows, the VxVM volume names must be 12 characters or fewer.
Otherwise, the backup fails.
- VxVM and VxFS_Checkpoint are the only snapshot methods in Snapshot Client that support the multi-volume file system (MVS) feature of VxFS 4.0.
- Since VxVM does not support fast mirror resynchronization on RAID 5 volumes, **VxVM** must not be used with VxVM volumes configured as RAID 5. If the **VxVM** snapshot method is selected for a RAID 5 volume, the backup fails.

Creating a snapshot mirror of the source

To use the **VxVM** snapshot method with VxVM volumes in the third-mirror (split-mirror) configuration, you must create a snapshot mirror. Use one of the following methods on the client.

To create a snapshot mirror of the source

- 1 In the Volume Manager Storage Administrator interface:
 - UNIX: select the source (primary) volume, right-click it, and select **Snapshot** from the pop-up menu. In the **Volume Snapshot** dialog, select **Enable FMR** (if available: see the following note) and click the **Snapstart** option.
 - Windows: select the source volume, right-click it, select **Snap**, and choose **Snapstart**.
For details, refer to your Veritas Volume Manager documentation.
- 2 Or for UNIX, enter the following commands:

```
/usr/sbin/vxassist -g disk_group snapstart volume_name
/usr/sbin/vxvol -g disk_group set fmr=on volume_name
```

where:

- *disk_group* is the Volume Manager disk group to which the volume belongs.
- *volume_name* is the name of the volume that is designated at the end of the source volume path (for example, *vol1* in */dev/vx/rdisk/dg/vol1*).
- *fmr=on* sets the Fast Mirror Resynchronization attribute, which resynchronizes the mirror with its primary volume. This attribute copies only the blocks that have changed, rather than performing a full resynchronization. Fast mirror resynchronization can dramatically reduce the time that is required to complete the backup.
Fast Mirror Resynchronization (FMR) is a separate product for Veritas Volume Manager.

- 3 With the Media Server or Third-Party Copy method, the disks that make up the disk group must meet certain requirements.

See [“Disk requirements for Media Server and Third-Party Copy methods”](#) on page 152.

About VxVM instant snapshots

Snapshot Client supports two additional kinds of snapshot volumes that are included in Volume Manager 4.0: full-sized instant snapshots and space-optimized instant snapshots. These types of volumes offer some advantages over traditional third-mirror snapshots, such as immediate availability and easier configuration and administration:

- Full-sized instant snapshot
This type of snapshot volume is a variation on the VxVM third-mirror volume snapshot model. It makes a snapshot volume available for access as soon as the snapshot plexes have been created. Like traditional third-mirror volumes, this volume after resynchronization can be moved into a separate disk group or turned into an independent volume.
- Space-optimized instant snapshot
This type of snapshot volume contains only the blocks that changed during the snapshot and uses a storage cache (cache object) to store them. The size of this cache may be configured when the snapshot is created. This volume can be created very quickly and uses a minimum of disk space. Note that it cannot be moved into a separate disk group or turned into an independent volume. Refer to *Veritas Volume Manager 4.0 Administrator's Guide* for complete descriptions of instant snapshot volumes and for configuration assistance.

About NetBackup snapshot methods

The instant snapshot feature of VxVM 4.0 is supported by NetBackup's VxVM, FlashSnap, and VVR methods. Note the following:

- FlashSnap supports VxVM full-sized instant snapshots, but not space-optimized snapshot. Additionally, FlashSnap supports VxVM volumes in a shared disk group.
- For alternate client backup, only the VVR method supports space-optimized snapshots.

Apart from configuring the VxVM volumes and selecting VxVM, FlashSnap, or VVR as the NetBackup snapshot method, no special parameters in NetBackup are required.

Space-Optimized instant snapshots

To use the space-optimized snapshot feature of VxVM, you must create a cache object for the snapshot.

About FlashSnap

FlashSnap uses the Persistent FastResync and Disk Group Split and Join features of Veritas Volume Manager (VxVM).

The FlashSnap snapshot method can be used for alternate client backups only, in the split mirror configuration.

See [“Alternate client backup split mirror examples”](#) on page 22.

FlashSnap supports VxVM full-sized instant snapshots, but not space-optimized snapshot. Additionally, FlashSnap supports VxVM volumes in a shared disk group. For support configurations, please refer to the *NetBackup 7.x Snapshot Client Compatibility List*.

Testing volumes for FlashSnap

Before you run an alternate client backup with the FlashSnap snapshot method, test your FlashSnap volume configuration as described in this section. You must ensure that the disk(s) containing the volume can be moved (deported and imported) to the alternate client without errors.

For instance, the backup disks cannot be split off into a new disk group and deported, if the disks containing the backup volumes contain some part of a volume that has another part on a disk that is not in the backup disks. Deporting a disk group means disabling access to that disk group.

See the *Volume Manager Administrator's Guide* for more information on deporting disk groups.

The following steps are described in more detail in the *Veritas FlashSnap Point-In-Time Copy Solutions Administrator's Guide*.

To test volumes for FlashSnap on UNIX

1 On the primary host:

- Add a DCO log to the volume:

```
vxassist -g diskgroup addlog volume logtype=dcg
```

- Enable FastResync on the volume:

```
vxvol -g diskgroup set fastresync=on volume
```

- Create a new snapshot mirror:

```
vxassist -g diskgroup -b snapstart volume
```

- Create a snapshot volume from the primary volume:

For disk groups that were created with VxVM 5.x:

```
vxassist -g diskgroup snapshot volume snap_volume
```

Note: Choose a volume name that is fewer than 15 characters. Volume names are prefixed or suffixed with characters, which increases the volume name. Lengthy volume names can cause snapshot jobs to fail.

For disk groups that were created with VxVM 4.x:

```
vxsnap -g diskgroup snapshot volume snap_volume
```

For disk groups that were created with VxVM 3.x:

```
vxassist -g diskgroup snapshot volume snap_volume
```

- Move the disks containing the snapshot volume to a separate (split) disk group:

```
vxdg split diskgroup split_diskgroup snap_volume
```

If the volume has not been properly configured, you may see an error similar to the following:

```
host-name# vxdg split lhdvvr lhdvvr_split SNAP-emc_concat
vxvm:vxdg: ERROR: vxdg split lhdvvr lhdvvr_split failed
vxvm:vxdg: ERROR: emc_dis05 : Disk moving, but not all
subdisks on it
```

- Re-examine the layout of the disks and the volumes that are assigned to them, and reassign the unwanted volumes to other disks as needed. Consult the *Veritas FlashSnap Point-In-Time Copy Solutions Administrator's Guide* for examples of disk groups that can and cannot be split.

- Deport the split disk group:

```
vxdg deport split_diskgroup
```

2 On the secondary host:

- Import the disk group that was deported from the primary:

```
vxdg import split_diskgroup
```

- Enable the imported volume:

```
vxrecover -g split_diskgroup -m snap_volume
```

- Start the volume:

```
vxvol -g split_diskgroup start snap_volume
```

If these commands execute successfully, the volume setup is correct.

3 After this test, you must re-establish the original configuration to what it was before you tested the volumes.

- Deport the disk group on the alternate client.
- Import the disk group on the primary client.
- Recover and join the original volume group.

See [“Identifying and removing a left-over snapshot”](#) on page 188.

To test volumes for FlashSnap on Windows

1 On the primary host:

- If not already done, create a snapshot mirror:

```
vxassist snapstart \Device\HarddiskDmVolumes\diskgroup\volume
```

- Create a snapshot volume from the primary volume:

```
vxassist snapshot \Device\HarddiskDmVolumes\diskgroup\volume
DrivePath=C:\Temp\Mount SNAP-Volume
```

- Move the disks containing the snapshot volume to a separate (split) disk group.

The disk group is also deported after this command completes:

```
vxdbg -g DskGrp -n SPLIT-DskGrp split
\Device\HarddiskDmVolumes\diskgroup\snap_volume
```

2 On the secondary host:

- Rescan to make the deported disk group visible on the secondary host:

```
vxassist rescan
```

- Import the disk group that was deported from the primary:

```
vxdbg -g split_diskgroup import
```

- Assign the snapshot volume to an empty NTFS directory.

This example uses C:\Temp\Mount.

```
vxassist assign \Device\HarddiskDmVolumes\split_diskgroup
\isnap_volume DrivePath=C:\Temp\Mount
```

If these commands execute successfully, the volume setup is correct.

About VVR

The VVR snapshot method (for UNIX clients only) relies on the Veritas Volume Replicator, which is a licensed component of VxVM. The Volume Replicator maintains a consistent copy of data at a remote site. Volume Replicator is described in the *Veritas Volume Replicator Administrator's Guide*.

The VVR snapshot method can be used for alternate client backups only, in the data replication configuration.

See [“Alternate client backup through data replication example \(UNIX only\)”](#) on page 25.

VVR makes use of the VxVM remote replication feature. The backup processing is done by the alternate client at the replication site, not by the primary host or client.

VVR supports VxVM instant snapshots.

See [“About VxVM instant snapshots”](#) on page 102.

VVR volume replication configuration

Before you do a replication backup with VVR, make sure to configure the Volume Replicator as explained in the *Volume Replicator Administrator's Guide*.

VVR name registration

Inband Control (IBC) messages are used to exchange control information between primary and secondary hosts. A name has to be registered at both the primary and secondary hosts for each replicated volume group before IBC messaging can be used. The VVR snapshot method assumes that the application name is APP_NBU_VVR. To avoid an initial backup failure, you should register that name.

See [“Testing the replication setup with VVR”](#) on page 107.

If APP_NBU_VVR is not registered, NetBackup registers the name when the first backup is attempted, but the backup fails. Subsequent backups, however, succeed.

Primary and secondary disk group and volume names for VVR

For the VVR snapshot method, the disk group and volume must have the same name on both the primary host and secondary host. If the names are different, the VVR backup fails.

Testing the replication setup with VVR

Before you run an alternate client backup with VVR, test your replication setup as follows.

To test your replication setup with VVR

- 1 On both primary host and secondary host, register the APP_NBU_VVR name:

```
vxibc -g diskgroup register APP_NBU_VVR replicated_group
```

This command must be executed twice, once on each host.

- 2 On the primary host, send an IBC message to the secondary host:

```
vxibc -g diskgroup send APP_NBU_VVR replicated_group  
replication_link
```

- 3 On the secondary host, receive the IBC message from the primary host:

```
vxibc -g diskgroup -Rl0 receive APP_NBU_VVR replicated_group
```

- 4 On the secondary host, restart replication:

```
vxibc -g diskgroup unfreeze APP_NBU_VVR replicated_group
```

If these commands execute successfully, the replication setup is correct.

About NAS_Snapshot

NetBackup can make point-in-time snapshots of data on NAS (NDMP) hosts using the NDMP V4 snapshot extension. The snapshot is stored on the same device that contains the NAS client data. From the snapshot, you can restore individual files or roll back a file system or volume by means of the Instant Recovery feature.

Note: NetBackup for NDMP software is required on the server, and the NAS vendor must support the NDMP V4 snapshot extension.

You can control snapshot deletion by means of the **Maximum Snapshots (Instant Recovery Only)** parameter. This parameter is specified on the **Snapshot Options** dialog of the policy.

For detailed information about NAS snapshots, setting up a policy for NAS snapshots and format of NAS snapshot name, check the 'Network Attached Storage (NAS) snapshot configuration' chapter of this guide.

See [“Means of controlling snapshots”](#) on page 87.

About VSS

VSS uses the Volume Shadow Copy Service of Microsoft Windows and supports Instant Recovery. VSS is for local backup or alternate client backup.

For the most up-to-date list of Windows operating systems and disk arrays supported by this method, see the *NetBackup 7.x Snapshot Client Compatibility List* document available on the Veritas support site:

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

For alternate client backup, the client data must reside on either a disk array such as EMC, HP, or Hitachi with snapshot capability, or a Veritas Storage Foundation for Windows 4.1 or later volume with snapshots enabled. VSS supports file system backup of a disk partition (such as E:\) and backup of databases.

Note: VSS-based snapshot methods offer a general interface to Windows Shadow Copy Services. VSS selects the actual snapshot method depending on which snapshot provider is configured on the client. For instance, if the data resides on an EMC CLARiiON array and the array administrator configured the array and its snapshot capability, the Volume Shadow Copy Service selects the appropriate CLARiiON VSS hardware provider to take the snapshot.

For configuration assistance, refer to your Microsoft documentation.

Disk array reconfiguration steps

The following preconfiguration steps may be required, depending on the make and model of your disk array:

- For arrays in the EMC CLARiiON and DMX series, and for the HP EVA series, see the appropriate topic for your array.
- For backup of a disk array using the Windows VSS snapshot method with Instant Recovery, be sure to configure NetBackup disk array credentials (if required by the array) before you run the backup. A Point in Time Rollback fails if NetBackup did not have credentials to access the array during the backup.
See the appropriate credentials topic for your disk array and snapshot method.
- For Hitachi arrays that are set up for mirror-based backup, see Hitachi and HP arrays in the *NetBackup Snapshot Client Configuration* document. This document may be accessed from the following location:
<http://www.veritas.com/docs/000081320>

Configuration and testing of volumes with VSS

Before you run an alternate client backup with VSS, configure and test volumes.

See “[Testing volumes for FlashSnap](#)” on page 103.

Notes and restrictions on VSS

The following notes apply to local backup and to alternate client backup:

- Supports backup of Storage Foundation for Windows 4.1 or later logical volumes OR backup of file systems on supported disk arrays.
- Supports backup of Windows NTFS file systems on a disk partition, and backup of data in a database. The policy type can be MS-Exchange-Server or MS-Windows.
- Does not support the backup of Windows system-protected files (the System State, such as the Registry and Active Directory). If the volume containing the

data to back up includes Windows system files, that volume cannot be backed up with the VSS snapshot method.

- Does not support the backup of Windows system database files (such as RSM Database and Terminal Services Database).

Support for Cluster Volume Manager Environments (CVM)

This chapter includes the following topics:

- [About support for CVM environments](#)
- [Note on NetBackup and CVM](#)
- [About enabling VxVM or FlashSnap snapshots in a CVM environment](#)
- [About enabling the NetBackup client to execute VxVM commands on the CVM master node](#)

About support for CVM environments

The cluster functionality of the Veritas Volume Manager (VxVM) called CVM allows the nodes in a Veritas Cluster Server (VCS) environment to simultaneously access and manage disks under Veritas Volume Manager control.

The supported snapshot methods are VxVM, FlashSnap, and VxFS_Checkpoint.

Note the following:

- For the FlashSnap and VxVM snapshot methods, snapshots can be created on a client that is on a slave node of the CVM cluster.
- Instant Recovery point-in-time rollback can be performed even if the virtual name of the client to be restored has failed over to a different node. NetBackup obtains the required restore information from the master server.

- For the FlashSnap and VxVM snapshot methods, mirror synchronization of multiple volumes completes faster.
- For the FlashSnap snapshot method, the alternate client can be a CVM node.

If the alternate client is a CVM node and the

`/usr/opensv/netbackup/NB_SNC_ALLOW_SNAP_DG` touch file exists on that node, the snapshot disk group is imported as a shared disk group. Otherwise by default the snapshot disk group is imported as a private disk group.

For details on the cluster capabilities in the Volume Manager, refer to the *Veritas Volume Manager Administrator's Guide*.

Note that the cluster functionality of Veritas Volume Manager requires a separate license.

This chapter does not discuss Veritas Storage Foundation Cluster File System (SFCFS) or cluster management software such as Veritas Cluster Server (VCS). Such products are separately licensed, and are not included with Veritas Volume Manager. See the documentation that is provided with those products for more information about them.

Note on NetBackup and CVM

When adding clients to the NetBackup policy Client list, use the application's virtual name in the cluster, not an actual node name. The virtual name allows the same host to be backed up by the policy after it has failed over to another node. Otherwise, the backup attempts to access a node that is no longer accessible.

About enabling VxVM or FlashSnap snapshots in a CVM environment

To create a FlashSnap or VxVM snapshot on a client that is on a slave node of the cluster, enable the NetBackup client to execute commands on the CVM master node.

See [“About enabling the NetBackup client to execute VxVM commands on the CVM master node”](#) on page 113.

About enabling the NetBackup client to execute VxVM commands on the CVM master node

To back up a VxVM volume that is in a shared disk group on a CVM slave node, certain VxVM commands may have to be executed remotely on the CVM master node. Therefore, you must enable the NetBackup client to execute the commands on any node. (This requirement applies to the FlashSnap or VxVM snapshot methods only.)

For further instructions, see Enabling the NetBackup client to execute VxVM commands on the CVM master node in the following Veritas tech note:

<http://www.veritas.com/docs/000081320>

Configuration of snapshot methods for disk arrays

This chapter includes the following topics:

- [OS-specific configuration tasks](#)
- [About IBM DS6000 and DS8000 arrays](#)
- [About IBM DS4000 array](#)
- [About Hitachi SMS/WMS/AMS, USP/NSC, USP-V/VM](#)
- [About HP-XP arrays](#)
- [About array troubleshooting](#)

OS-specific configuration tasks

This topic describes the configuration tasks that are related to the NetBackup client host operating system.

About dynamic multi-pathing

In a SAN fabric, it is advantageous (for redundancy reasons) for multiple pathways to exist from a host computer to a device on a disk array. The purpose of dynamic multi-pathing software is the following:

- Manage these pathways so that only one is in use at a time
- Switch to another pathway if the one in use fails

Snapshot Client supports EMC PowerPath dynamic multi-pathing software on Solaris and Windows. Due to limitations that prevent the dynamic importing of devices, PowerPath is not supported on Linux.

For certain arrays, Snapshot Client works in a multipath configuration with or without PowerPath installed, because all visible device paths provide equal access to the array device. EMC Symmetrix and HP-EVA arrays behave in this manner. For other arrays, if multiple device paths exist to an array device, Snapshot Client can use only one of those paths. In this case, PowerPath must be used to ensure that Snapshot Client always uses the active device path. EMC CLARiiON behaves in this manner. If PowerPath is not in use and the disk array is EMC CLARiiON, the Fibre Channel switch must be configured so that only one path is zoned to the NetBackup client.

HBA configuration

The supported HBAs are Emulex and QLogic. The JNI HBA is not supported.

HBA persistent target bindings

Persistent target bindings must be configured for every HBA. Without persistent bindings, the array's target number cannot be guaranteed on the host. A persistent device target number ensures that snapshots appear at the same device location if a NetBackup client host restarts. Refer to your HBA vendor documentation for help configuring persistent bindings.

Note: Persistent target bindings are not needed if you use Leadville drivers on Solaris.

Verifying NetBackup client access, zoning, and LUN masking

You can use the `nbfirescan` command to verify that the NetBackup clients have access to the array devices and that the arrays are properly zoned and LUNs are LUN masked. Note that `nbfirescan` only displays LUNs that have actually been LUN masked to the host.

To verify NetBackup client access, zoning, and LUN masking

Enter the following on the client:

- UNIX

```
/usr/opensv/netbackup/bin/nbfirescan
```

This command queries the host's SCSI bus for all the SCSI (or Fibre) attached devices that are visible.

Example Linux output (wrapped to fit page):

DevicePath	Vendor	Product	ID	EnclosureId	DeviceId	[Ctl,Bus,Tgt,Lun]
/dev/sdb	DGC	RAID 5		APM00050602951	60:06:01:60:83:B0:11:00:4D:C4:8A:1D:	
35:EC:DA:11		[01,00,00,00]				
/dev/sdc	DGC	RAID 5		APM00050602951	60:06:01:60:83:B0:11:00:4C:C4:8A:1D:	
35:EC:DA:11		[01,00,00,01]				
/dev/sdd	DGC	RAID 5		APM00050602951	60:06:01:60:83:B0:11:00:4B:C4:8A:1D:	
35:EC:DA:11		[01,00,00,02]				
/dev/sde	DGC	RAID 5		APM00050602951	60:06:01:60:83:B0:11:00:4A:C4:8A:1D:	
35:EC:DA:11		[01,00,00,03]				
/dev/sdf	HP	HSV200		5000-1FE1-5007-0020	6005-08B4-0010-5F49-0000-5000-22F8-0000	
		[01,00,01,01]				
/dev/sdg	HP	HSV200		5000-1FE1-5007-0020	6005-08B4-0010-5F49-0000-5000-22FF-0000	
		[01,00,01,02]				

About IBM DS6000 and DS8000 arrays

The following sections include background information and configuration tasks for NetBackup Snapshot Client backups using IBM DS6000 and DS8000 arrays. These tasks must be completed before you run a backup.

IBM DS6000 and DS8000 software requirements

The following IBM software is required.

Table 9-1 Software that is required for IBM DS6000 and DS8000

Software	Where to install	Version
DSCLI	Default location	5.2.2.224 or higher

For instructions on installing the software, refer to your IBM documentation.

Preconfiguration for IBM arrays

No preconfiguration steps are required for IBM DS6000 and DS8000 arrays.

Configuring NetBackup to access the IBM DS6000 or DS8000 array

You must supply logon credentials that allow the NetBackup client to access the IBM array.

To configure NetBackup to access the IBM DS6000 or DS8000 array

- 1 In the NetBackup Administration Console, click the **Media and Device Management > Credentials > Disk Array Hosts** node.
- 2 Right-click in the **Disk Array Hosts** pane and select **New Disk Array Host**.
- 3 For the IBM DS6000 or DS8000 array, enter the name of the host management console (the system where the Storage Manager resides).
- 4 Select **IBM System Storage** in the **Disk Array Host Type** pull-down menu.
- 5 Enter the user name and password for the array.
- 6 Clear (uncheck) the **Connect using port number** box.

Configuring array access for NetBackup hosts not named in a policy

Certain NetBackup hosts that are not named as clients in a policy must be explicitly enabled to access array credentials. An example is a media server that is used for off-host backup processing but is not included in any policy's Clients list.

To configure array access for NetBackup hosts not named in a policy

- 1 In the NetBackup Administration Console, click **Host Properties > Master servers > double click name of master server > Properties > Credential Access**.
- 2 Click **Add** to enter the name of the client. Then click **OK**.

Configuring the IBM array for NetBackup

You must add each NetBackup client and alternate client to the IBM array and make array devices available to the clients. In brief, the steps are the following:

To configure the IBM array for NetBackup

- 1 On the IBM array, provide host name and port information for the NetBackup client.

Note the following:

- The nickname (on the IBM DS6000 or DS8000) can be the same as the NetBackup client name.

- In the IBM DS6000 or DS8000 Storage Manager interface, the host type for AIX may not be obvious. Select `IBM pSeries`.
 - As part of the host definition, select the WWPN of the NetBackup client. Your NetBackup client must be properly zoned on the SAN to allow communication between it and the array.
- 2 Repeat step 1 for each NetBackup client or alternate client that uses the array.
 - 3 Create a volume group and associate the volume group with the NetBackup host you have defined on the array. For details, refer to your IBM documentation.
 - 4 Create logical volumes (or logical drives) for the volume group. This step makes the volumes or drives visible to the NetBackup client. For details, refer to your IBM documentation.

Using DSCLI commands to obtain unique IBM identifiers

The NetBackup policy requires entry of the array's Unique ID. If your array administrator provided LUN numbers for the devices, you must convert those LUN numbers into unique IDs for entry in the NetBackup policy **Snapshot Resources** pane. You can obtain the LUN unique IDs in either of two ways, as described in this topic.

The LUN ID of the primary and snapshot (clone) logical volume can be found from the array by means of DSCLI commands or the IBM Storage Manager interface.

To use DSCLI commands to obtain unique IBM identifiers

- 1 Find the host connection and its corresponding volume group by entering the following:

```
lshostconnect -dev enclosure_ID
```

Example:

```
dscli> lshostconnect -dev IBM.1750-6866123
Date/Time: December 17, 2007 4:18:02 PM IST IBM DSCLI Version: 5.2.2.224 DS:
IBM.1750-6866123
Name          ID    WWPN                HostType      Profile                portgrp volgrpID
ESSIOport
=====
oigtsol05     0000 100000000C956A9B4  Sun           SUN - Solaris          0 V11
all
oigtaix03     0022 100000000C969F60E  pSeries       IBM pSeries - AIX      0 V46
all
oigtaix02     0023 100000000C94AA677  pSeries       IBM pSeries - AIX      0 V47
all
```

- 2 Find the volumes presented to this volume group and to the host:

```
showvolgrp -dev enclosure_ID
            volume_group
```

Example:

```
dscli> showvolgrp -dev IBM.1750-6866123 V47
Date/Time: December 17, 2007 4:21:01 PM IST IBM DSCLI Version: 5.2.2.224 DS:
IBM.1750-6866123
Name oigtaix02
ID    V47
Type SCSI Mask
Vols 0002 0003 0004 0005 0006 0007 0008 0009 0031
```

The values listed for `Vols` are the LUN ids.

- 3 Find out which device on the host corresponds to a given logical volume.

Enter the following:

```
/usr/opensv/netbackup/bin/nbfirescan
```

To use the IBM Storage Manager web interface to obtain the device identifiers

- 1 In the Storage Manager, click **Real-time manager > Manage hardware > Host systems**.

Host systems: Real-time

Refresh Last refresh: Monday, December 17, 2007 4:28:48 PM IST

Select storage complex
rattle

Select	Nickname	Type	Host Ports	Attachments	Attachment Types	Storage Units
☐	oigtai01	IBM pSeries, RS6000 and RS6000 SP servers (AD)(pSeries)	1	1		IBMD5
☐	oigtai02	IBM pSeries, RS6000 and RS6000 SP servers (AD)(pSeries)	1	1		IBMD5
☐	oigtai03	IBM pSeries, RS6000 and RS6000 SP servers (AD)(pSeries)	1	1		IBMD5
☐	oigtai06	AMD Servers - Linux RHEL (AMDLinuxRHEL)	1	1		IBMD5
☐	oigtai10	Intel-based servers (Linux RHEL)(LinuxRHEL)	1	1		IBMD5
☐	oigtai16	Intel-based servers (Linux RHEL)(LinuxRHEL)	1	1		IBMD5
☐	oigtai24	Intel-based servers (Linux RHEL)(LinuxRHEL)	2	2	FcSf	IBMD5
☐	oigtai00	Sun servers (Solaris) (Sun)	1	1		IBMD5
☐	oigtai04	Sun servers (Solaris) (Sun)	1	1		IBMD5
☐	oigtai05	Sun servers (Solaris) (Sun)	1	1		IBMD5

Page 2 of 4 | 2 Go Total: 31 Filtered: 10 Displayed: 0 Selected: 0

- 2 Click the host for which you need to find the volumes presented.
The volume groups that are associated with the host are displayed.
- 3 Click the volume group to get the list of the logical volumes that are configured in this volume group.
The **Number** column indicates the LUN ID.

Volumes - Open systems: Real-time

Select storage unit
IBMD5

Select secondary filter
Volume group

Select volume group
oigtai02

Refresh Last refresh: Monday, December 17, 2007 4:31:51 PM IST

Select	Nickname	Number	Status	Type	GB(2*30)	GB(10*9)	Blocks	RAID	Extent Pool	Volume Groups	Host Attach
☐	Flash01_source	0002	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	
☐	Flash02_source	0003	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	
☐	Flash03_source	0004	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	2
☐	Flash04_source	0005	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	
☐	Flash05_source	0006	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	2
☐	Flash01_target	0007	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	2
☐	Flash02_target	0008	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	
☐	Flash03_target	0009	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	2
☐	Flash03_s	0031	Normal	DS	1.0		2,097,152	RAID 5	Extent Pool 1	oigtai02	3

Page 1 of 1 Total: 9 Filtered: 9 Displayed: 9 Selected: 0

Configuring a NetBackup policy for IBM_DiskStorage_FlashCopy

For general help setting up a NetBackup policy in the NetBackup Administration Console, refer to the [NetBackup Administrator's Guide, Volume I](#).

The following procedure focuses on the IBM_DiskStorage_FlashCopy method and its parameters.

To configure a NetBackup policy for IBM_DiskStorage_FlashCopy

- 1 In the NetBackup Administration Console, click the **Perform snapshot backups** box on the policy **Attributes** tab.
- 2 Click the **Options** option to display the **Snapshot Options** dialog box.
- 3 In the **Snapshot method** pull-down list, select **IBM_DiskStorage_FlashCopy**.
- 4 If needed, set **Wait for flashcopy operation to complete** to 1.

By default (0), the backup does not wait for the FlashCopy operation to complete. If the backup begins before the FlashCopy operation completes, database performance on the client (such as Oracle) may be affected until the FlashCopy is complete. A setting of 1 means that NetBackup waits for the FlashCopy to complete before the backup begins, to avoid any performance issues on the client.

Note that a setting of 1 may cause the backup elapsed time to increase significantly. The primary (production) volume may not be accessible until the FlashCopy command completes.

- 5 On the **Snapshot Options** dialog box in the **Snapshot Resources** pane, click **Add**.

NetBackup uses the information in the **Snapshot Resources** pane to correctly rotate through the clone LUNs when performing a FlashCopy clone backup. After all clone LUNs have been used in a backup, NetBackup automatically determines which clone is the oldest. NetBackup expires the oldest clone so that it can be reused for the current backup. If that clone represents the only backup image copy, NetBackup also expires the backup image associated with the clone.

Note that the policy Backup Selections determine one or more source LUNs for which snapshots are taken for backup. For each source LUN that is specified in the policy **Backup Selections** list, you must provide the information detailed in the following steps.

See [“Using DSCSI commands to obtain unique IBM identifiers”](#) on page 118.

- 6 In the **Add Snapshot Resource** dialog box, enter the array's serial number in the **Array Serial Number** field.

- 7 Enter the unique ID for the source LUN in the **Source Device** field.
- 8 Enter the unique IDs for the clone LUNs in the **Snapshot Device(s)** field. To enter multiple IDs, place a semicolon between them.

Note the following:

- The clone LUNs must be unmasked to the client (or alternate client) before you start a backup.
- For Instant Recovery backups, the Snapshot Device(s) entries determine where and in what order the snapshots are retained.

For further reference on IBM arrays

The following IBM documents may be helpful:

- The IBM System Storage DS6000 Series: Copy Services
<http://www.redbooks.ibm.com/redpieces/abstracts/sg246782.html>
<http://www.redbooks.ibm.com/abstracts/sg246783.html>
- The IBM System Storage DS8000 Series: Copy Services
<http://www.redbooks.ibm.com/abstracts/sg246787.html>
<http://www.redbooks.ibm.com/abstracts/sg246788.html>

About IBM DS4000 array

The following sections include background information and configuration tasks for NetBackup Snapshot Client backups using IBM DS4000 arrays. These tasks must be completed before you run a backup.

[Table 9-2](#) describes the snapshot method for the IBM DS4000.

Table 9-2 New snapshot method for the IBM DS4000 disk array

Method name	Description
IBM_StorageManager_FlashCopy	For full-volume copy (clone) snapshots on the IBM DS4000 series of arrays (excluding 4100), with SMcli version 9.60.

Array preconfiguration tasks

Before you configure a NetBackup policy, make sure that the following tasks have been completed.

[Table 9-3](#) describes the tasks.

Table 9-3 Array preconfiguration tasks

Array administration tasks	Where described
Install the disk array and its software, including appropriate licenses.	See your array documentation. See “IBM 4000 software requirements” on page 123.
Install supported HBAs on the NetBackup primary client and alternate clients.	See your HBA documentation.
Zone the client HBAs through the Fibre Channel switch, so the array is visible to the primary client and to any alternate clients.	See your Fibre Channel documentation.
Install NetBackup, and array vendor snapshot management software, on the NetBackup primary client and any alternate clients.	See the appropriate installation documentation.
Create and configure the Access Logical Drive for the host connection at the array. Configure logical drives on the array and make them visible to the host.	See your array documentation.

IBM 4000 software requirements

The following IBM software is required.

Table 9-4 Software that is required for IBM 4000

Software	Where to install	Version
SMclient	Default location on NetBackup client: /opt/IBM_DS4000/	9.60 or higher
SMruntime	Default location on NetBackup client: /opt/IBM_DS4000/	9.16 or higher

For instructions on installing the software, refer to your IBM documentation.

Verifying NetBackup client access, zoning, and LUN masking

You can use the nbfirescan command to verify the following: that NetBackup clients have access to the array devices, and that the arrays are properly zoned and LUNs are unmasked. Note that nbfirescan only displays LUNs that have been unmasked and mapped to the host.

To verify NetBackup client access, zoning, and LUN masking

Enter the following on the client:

```
/usr/opensv/netbackup/bin/nbfireshcan
```

This command queries the host's SCSI bus for all the SCSI (or Fibre) attached devices that are visible.

Example output from an AIX host, for Hitachi and IBM arrays, followed by a description:

```
DevicePath  Vendor  Product ID  EnclosureId  DeviceId  [Ctl,Bus,Tgt,Lun]
-----
/dev/hdisk8  HITACHI  OPEN-V-CM    10266        241       [00,00,144640,00]
/dev/hdisk9  HITACHI  OPEN-V       10266        840       [00,00,144640,01]
/dev/hdisk45  IBM      1814         FASTT        IBM.1814-600A0B800042212
2000000004804132760:0A:0B:80:00:42:33:6E:00:00:16:EF:48:DC:3C:1F [00,00
,327936,01]
/dev/hdisk46  IBM      1814         FASTT        IBM.1814-600A0B800042212
2000000004804132760:0A:0B:80:00:42:33:6E:00:00:16:F7:48:DC:57:F3 [00,00
,327936,02]
/dev/hdisk43  IBM      1814         FASTT        IBM.1814-600A0B800042212
2000000004804132760:0A:0B:80:00:42:33:6E:00:00:14:A4:48:AA:52:87 [00,00
,327936,03]
/dev/rdisk/c2t6d11s6 HITACHI  DF600F      6484         48        [00,00,00,00]
/dev/rdisk/c2t6d10s6 HITACHI  DF600F      6484         46        [00,00,00,00]
/dev/rdisk/c2t10d3s6 HITACHI  OPEN-V -SUN  45027        18        [00,00,00,00]
/dev/rdisk/c2t10d0s6 HITACHI  OPEN-V-CM   45027        0         [00,00,00,00]
```

Note the following descriptions:

DevicePath	Represents the actual access point for the device as it exists on the client host.
EnclosureId	Unique for each physical disk array.
DeviceId	Unique for a physical disk or virtual disk in an enclosure. The EnclosureId/DeviceId pair constitutes a host-independent designation of a particular physical or virtual disk within a disk array.
Ctl,Bus,Tgt,LUN	Controller, bus, target, and LUN numbers are the elements that designate a particular physical or virtual disk from the perspective of the client host computer.

Configuring NetBackup to access the IBM DS4000 array

You must supply logon credentials that allow the NetBackup client to access the IBM array.

To configure NetBackup to access the IBM DS4000 array

- 1 In the NetBackup Administration Console, click the **Media and Device Management > Credentials > Disk Array Hosts** node.
- 2 Right-click in the **Disk Array Hosts** pane and select **New Disk Array Host**.
- 3 Enter a dummy (substitute) host name for the IBM 4000 array. Do not enter the actual host name of the array.
- 4 Select **IBM System Storage** in the **Disk Array Host Type** pull-down menu.
- 5 Enter a user name and the correct password for the array.

The actual user name is not required for disk array operations. You can enter a dummy user name. The password, however, must be valid.

- 6 Clear (uncheck) the **Connect using port number** box.

Configuring the IBM 4000 array for NetBackup

You must add each NetBackup client and alternate client to the IBM array and make array devices available to the clients. In brief, the steps are the following.

To configure the IBM 4000 array for NetBackup

- 1 On the IBM array, provide host name and port information for the NetBackup client. Note the following:
 - Define the host and host group.
A given host group must have a single host, and host name(s) must exist for every HBA port on the host.
 - The host name and host group name on the IBM 4000 must be the same as the NetBackup client name. If more than one host-HBA port entry exists, the host name can differ from the client name.
 - In the IBM Storage Manager interface, select AIX as the host type.
 - As part of the host definition, select the WWPN of the NetBackup client.
Your NetBackup client must be properly zoned on the SAN to allow communication between it and the array.
- 2 Repeat step 1 for each NetBackup client or alternate client that uses the array.

- 3 For every client and host group added, map an Access Logical Drive on LUN number 7 or 31.
- 4 Create logical drives and map them to the host group. This step makes the logical drives visible to the NetBackup client. For details, refer to your IBM documentation.

Configuring a NetBackup policy for IBM_StorageManager_FlashCopy

For general help setting up a NetBackup policy, refer to the [NetBackup Administrator's Guide, Volume I](#).

The following procedure focuses on the IBM_StorageManager_FlashCopy method and its parameters.

To configure a NetBackup policy for IBM_StorageManager_FlashCopy

- 1 In the NetBackup Administration Console, click the **Perform snapshot backups** box on the policy **Attributes** tab.
- 2 Click **Options** to display the **Snapshot Options** dialog box.
- 3 In the **Snapshot method** pull-down list, select **IBM_StorageManager_FlashCopy**.
- 4 You can set the following parameters:

Maximum snapshots (Instant Recovery only) parameter The maximum value is 4 snapshots per source device.

Repository % of Base (100 for Instant Recovery) Determines the size of the IBM repository logical drive as a percentage of the primary device (base logical drive). The size can range from 1% to 100%. The more write activity that occurs on the primary drive, the more space the repository logical drive requires.

If the size of the primary is 500 GB and you set this parameter to 30%, the repository drive is set to 150 GB (30% of 500).

For Instant Recovery backups, the percentage is set to 100% regardless of any value that is entered here. The default (0) does not mean 0%: it means that the array determines the size of the repository drive (usually 20%).

For more details about the repository logical drive, refer to the IBM System Storage DS4000 Series and Storage Manager document.

About Hitachi SMS/WMS/AMS, USP/NSC, USP-V/VM

The following sections include background information and configuration tasks for NetBackup Snapshot Client backups using Hitachi SMS/WMS/AMS, USP/NSC, and USP-V/VM series of arrays.

These tasks must be completed before you run a backup.

Hitachi array software requirements

The following Hitachi software is required.

Table 9-5 Software that is required for Hitachi arrays

Software	Where to install	Version
RAID Manager/LIB 64-bit	NetBackup client	01-12-03/04 or later. Note: To determine the RAID Manager version, run the following: /usr/lib/RMLIB/bin/whatrmver See “Determining the current RAID Manager version” on page 127.

For instructions on installing the software, refer to your Hitachi documentation.

Determining the current RAID Manager version

Use the following procedure.

To determine the current RAID Manager version

Enter the following:

```
/usr/lib/RMLIB/bin/whatrmver
```

Example output:

```
Model :RAID-Manager/LIB/Solaris  
Ver&Rev:01-12-03/04
```

Preconfiguration for Hitachi

Before running backups, you must configure the arrays and appropriate volumes. Refer to your Hitachi array documentation.

Pair status must be PSUS

After creating volume pairs, you must split each pair and leave the status of the pair at PSUS.

You can use the following CCI command:

```
pairsplit -g dg_name -d device_name -l
```

where *dg_name* and *device_name* are the names specified in the CCI configuration file for the primary device.

Configure command devices on NetBackup client and alternate client

The Hitachi command devices must be visible to the NetBackup client as well as to any alternate client. To configure command devices, refer to your Hitachi documentation.

About communication between NetBackup and the Hitachi array

You do not need to configure array credentials for the Hitachi array. All communication between NetBackup and the array is done by means of command devices. If multiple Hitachi arrays are connected to a NetBackup client, NetBackup sends the NetBackup command to the correct Hitachi array.

To determine if the command devices are visible, see the following topic:

See [“Determining if the Hitachi command devices are visible”](#) on page 128.

Determining if the Hitachi command devices are visible

Use the following procedure to determine if the Hitachi command devices are visible.

To determine if the Hitachi command devices are visible

Enter the following:

```
/usr/opensv/netbackup/bin/nbfirescan
```

Example output:

```
nbfirescan v7.0 - Copyright (c) 2016 Veritas Technologies LLC.
Rescanning devices.....Complete.
Device count: 1
Device count: 1

DevicePath Vendor Product ID EnclosureId DeviceId [Ctl,Bus,Tgt,Lun]
-----
/dev/sda    VMware Virtual disk - - [00,00,00,00]
```

The last line of the example shows a command device (DF600F-CM).

See [“About communication between NetBackup and the Hitachi array”](#) on page 128.

About configuring the Hitachi array for NetBackup

You must add each NetBackup client and alternate client to a host group on the Hitachi array.

Note the following:

- The name of the host group must match the host name or fully qualified domain name of the client as specified in the NetBackup policy. The name should not be more than 16 characters in length.
- For host entries in the host groups, specify the WWNN/WWPN of each host.

For legacy disk snapshot methods for hitachi, see the *NetBackup Snapshot Client Configuration* document. <http://www.veritas.com/docs/000081320>

Obtaining the Hitachi array serial number and the unique device identifiers

The NetBackup policy requires the Hitachi array's serial number and the unique IDs (device identifiers) for the source and clone LUNs. Use the following procedure to obtain that information.

To obtain the Hitachi array serial number and the unique device identifiers

Enter the following command:

```
/usr/opensv/netbackup/bin/nbfireshcan
```

Example output:

DevicePath	Vendor	Product ID	EnclosureId	DeviceId	[Ctl,Bus,Tgt,Lun]
/dev/rdisk/c2t6d15s	HITACHI	F600F	6484	53	[00,00,00,00]
/dev/rdisk/c2t6d14s6	HITACHI	F600F	6484	52	[00,00,00,00]

The Enclosure ID is the serial number and the Device ID is the array's device ID.

Configuring a NetBackup policy for Hitachi_ShadowImage or Hitachi_CopyOnWrite

For general help setting up a NetBackup policy in the NetBackup Administration Console, refer to the [NetBackup Administrator's Guide, Volume I](#).

The following procedure focuses on the Hitachi_ShadowImage and Hitachi_CopyOnWrite methods and their parameters.

Note: The term "clone LUNs," as used in this procedure, refers to the Hitachi_ShadowImage method. For the Hitachi_CopyOnWrite method, the term "clone LUNs" can be replaced with "snapshot LUNs."

To configure a NetBackup policy for Hitachi_ShadowImage or Hitachi_CopyOnWrite

- 1 In the NetBackup Administration Console, click the **Perform snapshot backups** box on the policy **Attributes** tab.
- 2 Click the **Options** option to display the **Snapshot Options** dialog box.
- 3 In the **Snapshot method** pull-down list, select **Hitachi_ShadowImage** or **Hitachi_CopyOnWrite**.

4 In the **Snapshot Resources** pane, click **Add**.

NetBackup uses the information in the **Snapshot Resources** pane to correctly rotate through the designated LUNs when performing a ShadowImage (or CopyOnWrite) backup. After all clone LUNs have been used in a backup, NetBackup automatically determines which clone is the oldest. NetBackup expires the oldest clone so that it can be reused for the current backup. If that clone represents the only backup image copy, NetBackup also expires the backup image associated with the clone.

Note that the policy's **Backup Selections** list must specify one or more source LUNs for which snapshots are created for the backup. For each source LUN in the policy **Backup Selections** list, you must supply a serial number and unique IDs.

Note: If the snapshot resources specified in the **Snapshot Resources** pane do not match with the LUN IDs specified through the backup selection entries then the NetBackup Hitachi storage array plug-in will discover the available target device and use it for snapshot.

See [“Obtaining the Hitachi array serial number and the unique device identifiers”](#) on page 129.

5 In the **Add Snapshot Resource** dialog box, enter the array's serial number in the **Array Serial #** field.

6 Enter the unique ID for the source LUN in the **Source Device** field.

The ID must be entered without leading zeroes. For example, if the LUN ID is 0110, enter 110 in the **Source Device** field.

7 Enter the unique IDs for the clone LUNs (for Hitachi_ShadowImage method) or the snapshot LUNs (for Hitachi_CopyOnWrite) in the **Snapshot Device(s)** field. To enter multiple IDs, place a semicolon between them.

The ID must be without leading zeroes. For example, if the LUN ID is 0111, enter 111 in the **Snapshot Device(s)** field.

Note the following:

- The LUNs must be unmasked to the client (or alternate client) before you start a backup.
- For Instant Recovery backups, the **Snapshot Device(s)** entries determine where and in what order the snapshots are retained.

About HP-XP arrays

The following sections include background information and configuration tasks for NetBackup Snapshot Client backups using HP-XP series of arrays.

These tasks must be completed before you run a backup.

HP-XP array software requirements

The following HP-XP software is required.

Table 9-6 Software that is required for HP-XP arrays

Software	Where to install	Version
RAID Manager/LIB 64-bit	NetBackup client	01-12-03/04 or later Note: To determine the RAID Manager version, run the following: <code>/usr/lib/RMLIB/bin/whatrmver</code> See “Determining the current RAID Manager version” on page 127.

For instructions on installing the software, refer to your HP-XP documentation.

Preconfiguration for HP-XP

Before running backups, you must configure the arrays and appropriate volumes. Refer to your HP-XP array documentation.

Pair status	After creating volume pairs, you must split each pair and leave the status of the pair at PSUS. You can use the following CCI command: <code>pairsplit -g dg_name -d device_name -l</code> where <i>dg_name</i> and <i>device_name</i> are the names specified in the CCI configuration file for the primary device.
Configure command devices on NetBackup client and alternate client	The HP-XP command devices must be visible to the NetBackup client as well as to any alternate client. To configure command devices, refer to your HP-XP documentation.

About communication between NetBackup and the HP-XP array

You do not need to configure array credentials for the HP-XP array. All communication between NetBackup and the array is done by means of command devices. If multiple HP-XP arrays are connected to a NetBackup client, NetBackup sends the NetBackup command to the correct HP-XP array.

To determine if the command devices are visible, see the following topic:

See [“Determining if the HP-XP command devices are visible”](#) on page 133.

Determining if the HP-XP command devices are visible

Use the following procedure to determine if the command devices are visible.

To determine if the command devices are visible

Enter the following:

```
/usr/opensv/netbackup/bin/nbfirescan
```

About configuring the HP-XP array for NetBackup

You must add each NetBackup client and alternate client to a host group on the HP-XP array.

Note the following:

- The name of the host group must match the host name or fully qualified domain name of the client as specified in the NetBackup policy. The name should not be more than 16 characters in length.
- For host entries in the host groups, specify the WWNN/WWPN of each host.

Obtaining the array serial number and unique HP-XP identifiers

The NetBackup policy requires the HP-XP array's serial number and the unique IDs (device identifiers) for the source and clone LUNs. Use the following procedure to obtain that information.

To obtain the array serial number and unique HP-XP identifiers

Enter the following command:

```
/usr/opensv/netbackup/bin/nbfirescan
```

Configuring a NetBackup policy for HP_XP_BusinessCopy and HP_XP_Snapshot

For general help setting up a NetBackup policy in the NetBackup Administration Console, refer to the [NetBackup Administrator's Guide, Volume I](#).

The following procedure focuses on the HP_XP_BusinessCopy and HP_XP_Snapshot methods and their parameters.

Note: The term "clone LUNs," as used in this procedure, refers to the HP_XP_BusinessCopy method. For the HP_XP_Snapshot method, the term "clone LUNs" can be replaced with "snapshot LUNs."

To configure a NetBackup policy for HP_XP_BusinessCopy and HP_XP_Snapshot

- 1 In the NetBackup Administration Console, click the **Perform snapshot backups** box on the policy **Attributes** tab.
- 2 Click the **Options** option to display the **Snapshot Options** dialog box.
- 3 In the **Snapshot method** pull-down list, select HP_XP_BusinessCopy or HP_XP_Snapshot.
- 4 In the **Snapshot Resources** pane, click **Add**.

NetBackup uses the information in the **Snapshot Resources** pane to correctly rotate through the designated LUNs when performing a BusinessCopy (or Snapshot) backup. After all clone LUNs have been used in a backup, NetBackup automatically determines which clone is the oldest. NetBackup expires the oldest clone so that it can be reused for the current backup. If that clone represents the only backup image copy, NetBackup also expires the backup image associated with the clone.

See ["Obtaining the array serial number and unique HP-XP identifiers"](#) on page 133.

- 5 In the **Add Snapshot Resource** dialog box, enter the array's serial number in the **Array Serial #** field.
- 6 Enter the unique ID for the source LUN in the **Source Device** field.

The ID must be entered without leading zeroes. For example, if the LUN ID is 0110, enter 110 in the **Source Device** field.
- 7 Enter the unique IDs for the clone LUNs (for HP_XP_BusinessCopy method) or the snapshot LUNs (for HP_XP_Snapshot) in the **Snapshot Device(s)** field. To enter multiple IDs, place a semicolon between them.

The ID must be without leading zeroes. For example, if the LUN ID is 0111, enter 111 in the **Snapshot Device(s)** field.

Note the following:

- The LUNs must be unmasked to the client (or alternate client) before you start a backup.
- For Instant Recovery backups, the **Snapshot Device(s)** entries determine where and in what order the snapshots are retained.

About array troubleshooting

These topics provide troubleshooting assistance for array-related issues.

Troubleshooting issues pertaining to all arrays

Note the following issues:

-
- Backups fail with the following message in the bpfis log:

```
snapshot services: snapshot method analysis failed: no combos  
generated: stack elements not capable of any split or quiesce.
```

This message can appear in the following cases:

- The source device or target devices were incorrectly specified in the policy.
- The wrong type of disk array method was selected. For example, if you selected an EMC CLARiiON method for an EMC Symmetrix device.

Troubleshooting Solaris issues

Backups fail on non-Leadville Solaris systems and the following message appears in the bpfis log:

```
devicefi: Failed to the initialize the import node  
"UDID##HP##HSV##5000-1FE1-5007-0020##6005-08B4-0010-5F49-0000-5  
000-901E-0000", device could not be found.
```

In this example, an HP-EVA snapshot was not found on the backup host. The `/kernel/drv/sd.conf` file probably has insufficient `lun=` entries. Add `lun=` entries for the HP-EVA target in `sd.conf` and restart the system. More information is available about LUN entries in `sd.conf`.

Troubleshooting host OS type specification

Some arrays (such as HP EVA) require that you specify the OS type of the host. If the OS type is not set, or is not set properly, unpredictable behavior may result.

Troubleshooting NetBackup and EMC CLARiiON arrays

Note the following issues:

Table 9-7 Issues with NetBackup and EMC CLARiiON arrays

Issue	Explanation/Recommended Action
Backups fail and the following message appears in the bpfis log: emccлариionfi: Unable to find SP that owns LUN 1. Verify that credentials have been supplied for the SP that owns LUN 1.	In this example, the message appears because LUN 1 is owned by a CLARiiON storage processor for which credentials have not been supplied. In a single path configuration, all LUNs must be owned by one storage processor and credentials must be supplied for at least that storage processor. It is okay to supply credentials for both storage processors because NetBackup automatically determines which set of credentials should be used. In a multipath configuration, credentials for both storage processors should be supplied because multipathing software may automatically change which storage processor owns the LUN.
Backups fail and the following message appears in the bpfis log: emccлариionfi: WARNING: Unable to import any login credentials for any appliances.	Credentials must be added for the CLARiiON array by means of the NetBackup Administration Console.
Backups fail and one or both of the following messages appear in the bpfis log: emccлариionfi: The host <i>hostname</i> was not found in any storage groups. To import a snapshot to host <i>hostname</i> , <i>hostname</i> must be in a Clariion storage group. emccлариionfi: LUN masking failed. Could not find a storage group containing the hostname [<i>hostname</i>].	NetBackup searches the CLARiiON's storage groups for the import host. (For a local backup, the import host is the host where the source device is mounted. For an off-host backup, the import host is the alternate client.) When the host is found, the snapshot device is assigned to that storage group, thus making it visible to the import host where the backup can proceed. If the import host is not in any storage groups, the backup fails.

Table 9-7 Issues with NetBackup and EMC CLARiiON arrays (*continued*)

Issue	Explanation/Recommended Action
Backups fail and the following message appears in the bpfis log: <pre>emccлариionfi: No more available HLU numbers in storage group. LUN <i>LUN number</i> cannot be LUN masked at this time</pre>	The device cannot be imported to the host, because the maximum number of devices from the array is already imported to this host. Expire any unneeded backup images.
EMC_CLARiiON_Snapview_Clone backups fail and the following message appears in the bpfis log: <pre>emccлариionfi: Could not find LUN <i>LUN number</i> in clonegroup <i>clonegroup name</i></pre>	The clone target device does not exist in the clonegroup belonging to the source device. Either correct the target list in the policy or use Navisphere to add the target device to the source device's clone group.
Both types of CLARiiON backups fail with the following in the bpfis log: <pre>emccлариionfi: CLIDATA: Error: snapview command failed emccлариionfi: CLIDATA: This version of Core Software does not support Snapview</pre>	These messages appear when the Snapview software is not installed on the CLARiiON array. Snapview must be installed on the array before CLARiiON clone or snapshot backups can succeed. Please see the array documentation or contact EMC for more information.
Backups fail and the following message appears in the bpfis log: <pre>execNAVISECCLI: CLI Command [<i>CLI command</i>] failed with error [<i>error number</i>]</pre>	NetBackup uses naviseccli to send commands to the CLARiiON array. If naviseccli encounters an error, it is captured and placed in the bpfis log. The lines immediately following the above line should contain the output from naviseccli that indicates why the command failed.
After a point-in-time rollback from a Windows VSS backup that was made with the EMC CLARiiON Snapview Clone snapshot provider, all clones are fractured (split from the primary)	As a best practice, avoid performing a point-in-time rollback from a Windows VSS backup that was made with the EMC CLARiiON Snapview Clone snapshot provider, if one of the clones is configured for the policy has not been used for an Instant Recovery backup. After a rollback, all the clones are placed in a "fractured" state. (Fractured clones are no longer synchronized with the primary.) As a result, any clone that had not already been used for a backup is no longer available for a future Instant Recovery backup. If you must perform a point-in-time rollback before all clones have been used for backups, note which clones are still synchronized before you do the rollback. After the rollback, you can manually resynchronize the clones.

Table 9-7 Issues with NetBackup and EMC CLARiiON arrays (*continued*)

Issue	Explanation/Recommended Action
Policy validation fails for Standard policy with the following message: The error message is <code>Incorrect snapshot method configuration or snapshot method not compatible for protecting backup selection entries</code>.	Policy validation for a Standard policy created with the <code>EMC_CLARiiON_Snapview_Snapshot</code> fails with error 4201. The policy validation fails when the CLI is installed at a location where NetBackup fails to identify it. The CLI must be installed in the <code>/sbin/naviseccli</code>. If the CLI is installed at another location, NetBackup fails to identify that location and policy validation fails. To fix the policy validation, add the following entry into the <code>emccлариionfi.conf</code> file: <pre>[CLI_TOOL_INFO] "FILEPATH_NAVISEC_EXE"="<code><NAVI CLI path></code>" "FILENAME"="<code><NAVI CLI tool file name></code>"</pre> For example, for the HP platform the path would be: <pre>/usr/openv/lib/vxfi/configfiles/emccлариionfi.conf</pre> The entry would be: <pre>[CLI_TOOL_INFO] "FILEPATH_NAVISEC_EXE"="/opt/Navisphere/bin" "FILENAME_NAVISEC_EXE"="naviseccli"</pre>

Troubleshooting NetBackup and EMC Symmetrix arrays

Note the following issues:

Table 9-8 Issues with NetBackup and EMC Symmetrix arrays

Issue	Explanation/Recommended Action
Point in time rollback fails and the following message appears in the <code>bpfis</code> log: <code>Invalid clone state, cannot restore from device-ID to device-ID, where the first device ID is the source and the second is the clone.</code>	

Table 9-8 Issues with NetBackup and EMC Symmetrix arrays (*continued*)

Issue	Explanation/Recommended Action
If all Save Device space is consumed on the Symmetrix, a backup with EMC_TimeFinder_Snap or EMC_TimeFinder_Clone fails with the following error in the bpfis log: An internal Snap or Clone error has occurred. Please see the symapi log file	Check the symapi log (often found at <code>/var/symapi/log</code> on UNIX) to determine the exact error. If the log indicates there is no Save Device space, add Save Devices to the Save Device pool on your Symmetrix array.
EMC_TimeFinder_Mirror backups fail and the following message appears in the bpfis log: emcsymfi: Invalid STD-BCV pair state	This message indicates that the STD-BCV pair is not in a state that allows the mirror to be created. Verify that the pair were fully synchronized before the backup attempt.
TimeFinder backups fail and a message similar to the following appears in the bpfis log: device 00A4 to director/port that is not accessible to the HBA port 210000e08b86b34f emcsymfi: 2 of 2 HBA WWNs assigned to director/port combinations for accessing device 00A4 are not accessible due to misconfiguration of the array. emcsymfi: Since there were no valid mappings the device import will fail. To resolve this issue you must reconfigure your array 000187910258 to bind device 00A4 to director ports accessible to host.	Device 00A4 has not been mapped to any Symmetrix director ports that are zoned to the host. All Symmetrix devices (source devices and target devices) must be mapped to an array director port that is zoned to the import host. For a local backup policy, the import host is the host where the source device is mounted. For an off-host backup policy, the import host is the alternate client.

Troubleshooting NetBackup and HP EVA arrays

Table 9-9 Issues with NetBackup and HP EVA arrays

Issue	Explanation/Recommended Action
Backups fail with the following warning message in the bpfis log: WARNING: No credentials found for HP HSV	Credentials must be added for the EVA array by means of the NetBackup Administration Console.

Table 9-9 Issues with NetBackup and HP EVA arrays (*continued*)

Issue	Explanation/Recommended Action
Snapshot job fails when the client has VxVM software installed, but the underlying disk in Snapshot Client backup is not configured on stack. The following error message is displayed: <code>client/server handshaking failed (26) " for HP_EVA_Snapclone FIM</code>	Uninstall the VxVM software from the client.

Troubleshooting IBM DS6000 and DS8000 arrays

This section provides explanations and recommended actions, as well as log entries that may help identify the problem.

Note the following important items:

- Before you start a backup, the snapshot device (clone) must be visible (unmasked) to the NetBackup client or alternate client.
- For backup and Instant Recovery rollback restore: you must supply logon credentials that allow the NetBackup client to access the IBM array. For NetBackup hosts not named in a policy (such as alternate clients), you must also configure NetBackup so that the host can access the credentials. Backups may appear to succeed, but the bprd log on the server contains messages similar to the following:

```
09:02:17.999 [4292.3092] <2> is_disk_client_configured:  
db_cred_allowed(host1.enterprise.com, 1) failed: 227  
09:02:17.999 [4292.3092] <2> read_text_file:  
is_disk_client_configured(host1.enterprise.com) failed: 227  
09:02:17.999 [4292.3092] <2> process_request: read_text_file  
failed - status = client is not validated to use the server  
(131)
```

Until credential access is enabled, a backup or a point-in-time rollback fails with NetBackup status 5.

See [“Configuring NetBackup to access the IBM DS6000 or DS8000 array”](#) on page 117.

See [“Configuring array access for NetBackup hosts not named in a policy”](#) on page 117.

Snapshot errors with IBM DS6000 and DS8000 arrays (NetBackup status code 156)

Note these possible explanations:

Table 9-10 Issues encountered with Snapshot (NetBackup status code 156)

Issue	Explanation/Recommended Action
The snapshot device (clone) is not visible (unmasked) to the NetBackup client or alternate client.	Make the clone device visible to the NetBackup client or alternate client before you retry the backup. Contact IBM technical support or refer to your IBM array documentation.
The snapshot device (clone) is also a source device in another device pair. The following message may appear in the <code>/usr/opensv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log: <pre>CMUN03041E mkflash: Copy Services operation failure: already a FlashCopy source</pre>	Reconfigure source and clone devices so that the clone required for this backup is not a source device for another clone. Contact IBM technical support or refer to your IBM array documentation.
The snapshot device (clone) and source device are not of equal size. The following message may appear in the <code>/usr/opensv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log: <pre>CMUN03049E mkflash: Copy Services operation failure: incompatible volumes</pre>	Reconfigure source and clone devices to be identical in size. Contact IBM technical support or refer to your IBM array documentation.
The source device is already recording enabled for FlashCopy. The following message may appear in the <code>/usr/opensv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log: <pre>CMUN03027E mkflash: FlashCopy operation failure: action prohibited by current FlashCopy state. Contact IBM technical support for assistance.</pre>	Verify whether the source device has a FlashCopy relationship with some device other than the snapshot device (clone) specified in the policy. If a FlashCopy relationship exists with some other device, delete the relationship and start the backup again.

Table 9-10 Issues encountered with Snapshot (NetBackup status code 156)
(continued)

Issue	Explanation/Recommended Action
The IBM FlashCopy license is not installed. The following message may appear in the <code>/usr/opensv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log: <pre>CMUN03035E mkflash: Copy Services operation failure: feature not installed.</pre>	Install the FlashCopy license on the storage subsystem. Contact IBM technical support or refer to your IBM array documentation.
The FlashCopy relationship is not recording enabled. The following message may appear in the <code>/usr/opensv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log: <pre>CMUN03027E resyncflash: FlashCopy operation failure: action prohibited by current FlashCopy state. Contact IBM technical support for assistance</pre>	■ Make sure a FlashCopy relationship exists for the device pair. ■ If the FlashCopy relationship is not recording enabled, remove the FlashCopy relationship and then re-run the backup.
A FlashCopy relationship does not exist. The following message may appear in the <code>/usr/opensv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log: <pre>CMUN03027E resyncflash: FlashCopy operation failure: action prohibited by current FlashCopy state. Contact IBM technical support for assistance.</pre> A resynchronize operation was attempted on a FlashCopy pair that does not exist.	Verify that a FlashCopy pair does not exist, and then re-execute the backup.

Table 9-10

Issues encountered with Snapshot (NetBackup status code 156)

(continued)

Issue	Explanation/Recommended Action
Incremental Copy is in progress. The following message may appear in the <code>/usr/openv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log: <pre>CMUN02498E resyncflash: The storage unit is busy.</pre> This message can appear for the following reasons: ■ A background copy is in progress for the given pair of devices on the array. ■ Some maintenance activity is currently in progress on the array.	Allow the background copy or maintenance activity to complete. Then re-run the backup.

Troubleshooting IBM4000 arrays

These sections provide explanations and recommended actions, as well as log entries that may help identify the problem.

Snapshot errors encountered with IBM4000 arrays (NetBackup status code 156)

The following table lists explanations and recommended actions for status code 156.

Table 9-11 Explanations and recommended actions for status code 156

Problem	Description and recommended action for status code 156
The array does not have enough free space.	FlashCopy logical drives are created under the same logical array as is the base or primary logical drive. The storage subsystem might have free space, but if the logical array has insufficient space, the FlashCopy operation fails. The following messages may appear in the /usr/openv/netbackup/logs/bpfis log: <pre>23:44:48.007 [655588] <2> onlfi_vfms_logf: INF - snapshot services: ibmdsf:Wed Mar 12 2008 23:44:48.007721 <Thread id - 1> FlashCopy could not be created. command [create FlashCopyLogicalDrive baseLogicalDrive="angela_javelin_4" userLabel="angela_javelin_4_flcp_4"];23:44:48.012 [655588] <2> onlfi_vfms_logf: INF - snapshot services: ibmdsf: FlashCopy creation failed for source volume angela_javelin_4 on storage unit ibmds4700.</pre> In addition, the following messages may appear in /usr/openv/netbackup/logs/bpfis/ibmdsf.log.date: <pre>Mon Mar 31 2008 14:25:23.036588 <Pid - 1065104 / Thread id - 1> FlashCopy could not be created. command [create FlashCopyLogicalDrive baseLogicalDrive="drive-claix11-1" userLabel="drive-claix11-1_flcp"];. Mon Mar 31 2008 14:25:23.037164 <Pid - 1065104 / Thread id - 1> OUTPUT=[Unable to create logical drive "drive-claix11-1_flcp" using the Create FlashCopy Logical Drive command at line 1. Error - The operation cannot complete because there is not enough space on the array. The command at line 1 that caused the error is: create FlashCopyLogicalDrive baseLogicalDrive="drive-claix11-1" userLabel="drive-claix11-1_flcp";</pre> Recommended action: Make sure that the array has enough space available for the snapshot.
The maximum number of FlashCopies (4) already exist on the array.	The following messages may appear in /usr/openv/netbackup/logs/bpfis/ibmdsf.log.date: <pre>Mon Mar 31 2008 14:25:23.036588 <Pid - 1065104 / Thread id - 1> FlashCopy could not be created. command [create FlashCopyLogicalDrive baseLogicalDrive="drive-claix11-1" userLabel="drive-claix11-1_flcp"];. Mon Mar 31 2008 14:25:23.037164 <Pid - 1065104 / Thread id - 1> OUTPUT=[Could not create a flashcopy logical drive using the Create FlashCopyLogicalDrive command at line 1. Error 129 - The operation cannot complete because the maximum number of flashcopy logical drives have been created for this base logical drive.</pre> Recommended action: Delete any FlashCopies that NetBackup did not create.
The Access Logical Drive is not mapped for the NetBackup client or alternate client at LUN 31 or 7.	On the IBM DS4000, the Access Logical Drive communicates with the storage subsystem. Any client that is connected to and needs to communicate with the storage subsystem should have an Access Logical Drive mapped to it. If an Access Logical Drive is not mapped to the client, the client is unable to communicate with the array. As a result, any NetBackup client operation involving the array fails. Recommended action: Create and map an Access Logical Drive. Contact IBM technical support or refer to your IBM array documentation.

Table 9-11 Explanations and recommended actions for status code 156
(continued)

Problem	Description and recommended action for status code 156
The DAR driver is not functional.	Recommended action: Make sure that the RDAC package is installed on the AIX host.

For further reference on IBM System Storage DS4000

The following IBM documents may be helpful:

- *IBM System Storage DS4000 Series and Storage Manager*
<http://www.redbooks.ibm.com/abstracts/sg247010.html>
- *IBM System Storage DS Storage Manager Version 10.30: Copy Services User's Guide*
<http://www.filibeto.org/unix/aix/lib/hardware/ds4800/copy-services-ug-gc53113600.pdf>
- *IBM System Storage DS4000 Storage Manager Version 9 Installation and Support Guide*

Troubleshooting Hitachi arrays

This section provides explanations and recommended actions, as well as log entries that may help identify the problem.

Note the following important items:

- RAID Manager version 01-12-03/04 or later is required.
See “[Determining the current RAID Manager version](#)” on page 127.
- Before you start a backup, the snapshot device (clone LUN or snapshot LUN) must be visible (unmasked) to the NetBackup client or alternate client.
- For the NetBackup policy Snapshot Resource configuration, specify device IDs in decimal and without leading zeros.
For example, if your source device ID is 0100 and the snapshot device ID is 0101, enter 100 and 101 in the **Snapshot Resources** dialog box.
See “[Configuring a NetBackup policy for Hitachi_ShadowImage or Hitachi_CopyOnWrite](#)” on page 130.

NetBackup policy validation fails with Hitachi arrays

Note these possible explanations:

Table 9-12 NetBackup policy validation fails

Issue	Explanation/Recommended Action
Look for the following error in the <code>/usr/opensv/netbackup/logs/bpfis/hitachi.log.<date></code> log: <pre>Library RMLIB init failed</pre>	Make sure that the RMLIB 64-bit library is installed. This requirement applies when you upgrade from a 6.5.x system (requires 32-bit RMLIB) to a 7.1 system, and when you install a fresh 7.1 system.
The Hitachi command device is not unmasked. See the sample log messages in the next row.	Refer to Hitachi documentation for creating and unmasking command devices.
The Hitachi command device is unmasked but is not visible to client, or the enclosure ID specified in the policy's Snapshot Resources is invalid. The <code>/usr/opensv/netbackup/logs/bpfis/hitachi.log.<date></code> log may contain messages similar to the following: <pre>Fri Mar 21 2008 16:26:46.431046 <Pid - 9477 / Thread id - 1> Entering Function delayedInit [110, providers/hitachi/hitachi_plugin.cpp] Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> <Device name="/dev/rdskc4t50060E801029F700d2s6" udid="UDID##HITACHI##HDS##75040816##3" bus="0" target="0" lun="0" vendor="HITACHI" product="DF600F-CM" /> Fri Mar 21 2008 16:26:49.174493 <Pid - 9477 / Thread id - 1> Exiting Function delayedInit [110, providers/hitachi/hitachi_plugin.cpp]</pre> If the <code>delayedInit</code> message does not include at least one entry for the enclosure ID that was entered in the policy's Snapshot Resources, the command device is not unmasked or is not visible to NetBackup client (host).	Make sure that the command device is recognized by the operating system and that the enclosure ID is entered correctly in the policy's Snapshot Resources. To determine if the command device is recognized by the operating system, try device discovery commands such as the following: <pre>devfsadm cfgadm -al</pre> A log message for the enclosure ID would include an entry such as the following: <pre>c3t50060E801029F700d28 <HITACHI-DF600F-CM-0000 cyl 52 alt 2 hd 50 sec 768></pre> which shows that the device is visible as <code>c3t50060E801029F700d28</code>.
A mismatch exists between the policy's snapshot method and the type of LUNs specified for the Snapshot Devices. For example, if you select the <code>Hitachi_ShadowImage</code> method but specify snapshot LUNs instead of clone LUNs for the Snapshot Devices, an error occurs. See the sample log messages in the next bullet.	Specify the correct snapshot method or snapshot devices.

Table 9-12 NetBackup policy validation fails (*continued*)

Issue	Explanation/Recommended Action
A disk pair was not created for the source device and snapshot device specified in the NetBackup policy's Snapshot Resources. The <code>/usr/openv/netbackup/logs/bpfi/hitachi.log.<date></code> log may contain messages similar to the following. If the snapshot method is Hitachi_CopyOnWrite: <pre>Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> HITACHI_FIM_SNAPSHOT not supported for 10</pre> If the snapshot method is Hitachi_ShadowImage: <pre>Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> HITACHI_FIM_SHADOW_IMAGE not supported for 10</pre>	Set up a disk pair (primary and secondary) for the source device and snapshot device that are specified in the policy's Snapshot Resources. Refer to the Hitachi documentation.
In the policy's Snapshot Resources, the device identifier for the source device or snapshot device is invalid. The <code>/usr/openv/netbackup/logs/bpfi/hitachi.log.<date></code> log may contain messages similar to the following: <pre>Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> getrminfo failed. Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> operation failed with error number <> with message <msg>'.</pre> The above message may indicate that a device ID in the policy's Snapshot Resources is incorrect or does not exist. For example, if the specified snapshot device ID does not exist: <pre>Mon May 12 2008 21:32:32.088876 <Pid - 8040 / Thread id - 1> getrminfo is called for '9999'. Mon May 12 2008 21:32:32.089736 <Pid - 8040 / Thread id - 1> getrminfo failed. Mon May 12 2008 21:32:32.090003 <Pid - 8040 / Thread id - 1> operation failed with error number '-1' with message '[EL_CMDRJE] An order of the control command rejected.'.</pre>	Recommended action: Make sure that the identifiers are correctly entered in the policy's Snapshot Resources. Specify source and snapshot IDs without leading zeros. See "Configuring a NetBackup policy for Hitachi_ShadowImage or Hitachi_CopyOnWrite" on page 130.

Table 9-12 NetBackup policy validation fails (continued)

Issue	Explanation/Recommended Action
The RAID Manager library <code>libsrvrm.so</code> software is not installed in the <code>/usr/lib/</code> directory.	Recommended action: Install the RAID Manager package in <code>/usr/lib/</code> . See the Hitachi documentation.
The installed version of RAID Manager library <code>libsrvrm.so</code> is not supported.	Recommended action: Look for the <code>Library RMLIB version</code> message in the <code>/usr/openv/netbackup/logs/bpfis/hitachi.log.<date></code> log. See “Determining the current RAID Manager version” on page 127.

Backup or restore or image expiration fails with Hitachi arrays

The following table lists possible explanations:

Table 9-13 Backup or restore or image expiration fails

Issue	Explanation/Recommended Action
	These messages indicate that the instance number is not available. A problem may exist with the instance number management logic. - Remove the following file. <code>/usr/opensv/lib/vxfi/cachefiles/hitachi/ UDID##HITACHI##HDS##enclosure id##*</code> The <i>enclosure id</i> is the array's serial number that is specified in the policy's Snapshot Resources. - Gather the appropriate Hitachi logs and contact Veritas customer support for NetBackup.

Table 9-13 Backup or restore or image expiration fails (continued)

Issue	Explanation/Recommended Action
A problem occurred that involved the Hitachi instance number. For example, the Hitachi snapshot provider did not receive the instance number for the command device. The instance number is needed to connect to the array. The <pre>/usr/openv/netbackup/logs/bpfis/hitachi.log.<date></pre> log contains the following message: <pre>Couldn't get instance no failed with message</pre> The log may contain the following additional messages: <pre>Fri Mar 21 2008 16:26:49.818233 <Pid - 9477 / Thread id - 1> Entering Function attachCmd [156, providers/hitachi/hitachi_rmlibintf.cpp] Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> attachcmddev is called with cmd dev /dev/dsk/clt0d0s2 and instance number 0. Fri Mar 21 2008 16:26:49.818308 <Pid - 9477 / Thread id - 1> Exiting Function attachCmd [156, providers/hitachi/hitachi_rmlibintf.cpp]</pre> The attachcmddev message should list the Hitachi command device (for the enclosure ID that was specified in the policy) and the instance number. If the Hitachi command device is not included in the message, then the instance number was not received. A limited number of instance numbers are allowed per command device. If the maximum number of processes is using the same command device, no more instance numbers are available. This situation may indicate a problem with the instance number management logic of the Hitachi provider. Note also the following potential messages: <pre>Fri Mar 21 2008 16:26:49.818233 <Pid - 9477 / Thread id - 1> Entering Function attachCmd [156, providers/hitachi/hitachi_rmlibintf.cpp] Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> Couldn't get instance no failed with message '%s'. Fri Mar 21 2008</pre>	

Table 9-13 Backup or restore or image expiration fails (continued)

Issue	Explanation/Recommended Action
16:26:49.818308 <Pid - 9477 / Thread id - 1> Exiting Function attachCmd [156, providers/hitachi/hitachi_rmlibintf.cpp] Another example message is the following: Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> Couldn't get instance no failed with message 'Instance No Exhausted, couldn't reclaim giving up'	
The default array controller of the source device is not the same as the controller of the snapshot device. Use the Storage Navigator interface to verify.	Recommended action: Make sure that the clone (or snapshot) device has the same default controller as the source device. See the Hitachi documentation.

Notes on Media Server and Third-Party Copy methods

This chapter includes the following topics:

- [Disk requirements for Media Server and Third-Party Copy methods](#)
- [Directives for Media Server and Third-Party Copy methods](#)
- [Storage units for Media Server and Third-Party Copy methods](#)
- [Preventing multiplexing on a third-party copy backup](#)
- [Raw partition backups](#)
- [Increasing the client read timeout for all clients](#)
- [Further information on off-host data mover backups](#)

Disk requirements for Media Server and Third-Party Copy methods

For the NetBackup Media Server or Third-Party Copy Device backup method, the client's data must be on one or more disks that meet the following criteria.

- The disk must be either a SCSI or Fibre Channel device.
- The disk must be visible to both the NetBackup client and to the NetBackup media server. The disk must be connected through a Fibre Channel SAN or through a disk array that has dual port SCSI connections.

- The disk must be able to return its SCSI serial number in response to a serial-number inquiry (serialization). Or, the disk must support SCSI Inquiry Page Code 83.

Directives for Media Server and Third-Party Copy methods

The policy's **Backup Selections** must not contain the ALL_LOCAL_DRIVES entry (except for VMware).

Storage units for Media Server and Third-Party Copy methods

Note the following:

- Any_available is not supported for NetBackup Media Server and Third-Party Copy Device backup methods.
- Disk storage units are not supported for the Third-Party Copy Device method.

Preventing multiplexing on a third-party copy backup

The Third-Party Copy Device backup method is incompatible with multiplexing (the writing of two or more concurrent backup jobs to the same storage device). To prevent multiplexing on a third-party copy backup, you must set **Maximum multiplexing per drive** to 1. Make this setting on the **Add New Storage Unit** or **Change Storage Unit** dialog.

To prevent multiplexing on a third-party copy backup

On the **Add New Storage Unit** or **Change Storage Unit** dialog, set **Maximum multiplexing per drive** to 1.

Raw partition backups

For the NetBackup Media Server or Third-Party Copy method, do not specify a block device as the raw partition to back up. For these two backup methods, NetBackup does not support block devices. Instead, specify the raw partition as a character device.

Examples:

```
Solaris: /dev/rdisk/clt3d0s3  
HP: /dev/rdisk/clt0d0
```

Increasing the client read timeout for all clients

For the NetBackup Media Server method, it may be necessary to increase the client read timeout value. In some environments, NetBackup may require more read time than the default value allows. If the client read timeout is insufficient, the backup may fail with status 13, `file read failed`.

To increase the client read timeout for all clients

- 1 In the NetBackup Administration Console, go to **Host Properties > Master Servers > *double click the master server* > Properties > Timeouts**.
- 2 Increase the Client read timeout.
- 3 Retry the backup.

Further information on off-host data mover backups

See the off-host data mover backups section in the *NetBackup Snapshot Client Configuration* document:

<http://www.veritas.com/docs/000081320>

Backup and restore procedures

This chapter includes the following topics:

- [About performing a backup](#)
- [About performing a restore](#)
- [Instant Recovery restore features](#)
- [Notes for restoring individual files from an Instant Recovery snapshot](#)
- [About configurations for restore](#)
- [About restoring from a disk snapshot](#)

About performing a backup

The following types of backups can be used with Snapshot Client policies.

Automatic backup

The most convenient way to back up client data is to configure a policy and then set up schedules for automatic, unattended backups. To use NetBackup Snapshot Client, you must enable snapshot backup as described in the appropriate configuration chapter of this guide. To add new schedules or change existing schedules for automatic backups, you can follow the guidelines in the [NetBackup Administrator's Guide, Volume I](#).

Manual backup	The administrator can use the NetBackup Administration interface on the master server to execute a backup for a policy. To use NetBackup Snapshot Client, you must enable snapshot backup as described in the appropriate configuration chapter of this guide. See the NetBackup Administrator's Guide, Volume I, for instructions on doing manual backups.
User-directed backup and archive	From a NetBackup client, the user can execute a Snapshot Client backup. The NetBackup administrator must configure an appropriate snapshot policy with schedule. See the NetBackup Backup, Archive, and Restore Getting Started Guide for instructions on doing user-directed backups and archives.

About performing a restore

You can use the Backup, Archive, and Restore interface to restore individual files or directories, or a volume or raw partition. See the [NetBackup Backup, Archive, and Restore Getting Started Guide](#) for instructions on performing the restore. The following topics include restore notes and procedures unique to certain components of Snapshot Client.

Instant Recovery restore features

You can restore files from an Instant Recovery backup in the same way as from a normal backup. Restore procedures are described in the *NetBackup Backup, Archive, and Restore Getting Started Guide*.

In addition, note the features unique to Instant Recovery may require special instructions.

About Instant Recovery: point in time rollback

You can also restore a snapshot of an entire file system or volume with minimal I/O. This type of restore is called point in time rollback. All the data in the snapshot is restored; single file restore is not available in a rollback.

Notes on rollback

Note the following.

Warning: Rollback deletes all files that were created after the creation-date of the snapshot that you restore. Rollback returns a file system or volume to a given point in time. Any data changes or snapshots that were made after that time are lost.

Also, if there are multiple logical volumes on a single disk or volume group and if you perform a **Point in Time Rollback** of a specific logical volume, the entire disk or volume group is restored to the point in time.

Performing snapshot rollback

The following procedure requires root access (UNIX) or Administrator privilege (Windows).

To perform snapshot rollback (UNIX)

- 1 Start the Backup, Archive, and Restore interface.

```
/usr/opensv/netbackup/bin/jbpaSA &
```


- 2 Click the **Restore Files** tab.
- 3 Click **Actions > Specify NetBackup Machines** to specify the server, source client, policy type, and destination client.
- 4 For the Restore Type, select **Point in Time Rollback**.

The **Browse directory** field is grayed out, with root (/) as default.

Instant Recovery backups are displayed in the **Backup History** window, for all dates (you cannot set a range).

- 5 Select an image from the list and click **OK**.

The image contents are displayed in the **Directory Structure** pane of the **Restore Files** tab.




You can select root level or mount points (file systems or volumes), but not folders or files at a lower level.

- 6 In the **Directory Structure** list, click the check box next to the root node or a mount point beneath root.

You can select a file system or volume, but not lower-level components.

- 7 Click the **Restore** option.



If the volume that you have selected, belongs to a storage array consistency group or volume set, you can see a list of volumes that are restored along with the selected volume, as they belong to the same array volume set or consistency group. In the open dialog,

- To recover all the listed volumes, click **OK**.
- To cancel, and select another restore method for single volume recovery, click **Cancel**.

For more details refer to NetBackup status code: 4311, in the *NetBackup Status Codes Reference Guide*

The only available destination option is **Restore everything to its original location**.

- 8 For file systems, you can choose to skip file verification by placing a check in the **Skip verification and force rollback** option.

Warning: Click **Skip verification and force rollback** only if you are sure that you want to replace all the files in the original location with the snapshot. Rollback deletes all files that were created after the creation-date of the snapshot that you restore.

If **Skip verification and force rollback** is not selected, NetBackup performs checks on the file system.


See [“Notes on rollback”](#) on page 156.

If the checks do not pass, the rollback aborts and the **Task Progress** tab states that rollback could not be performed because file verification failed.

- 9 For volumes belonging to consistency group or volume set, select **Force rollback even if it destroys the consistency group's state on the storage array** option, to ignore any mismatch in consistency group configuration on the storage arrays during the restore. NetBackup completes the restore with warning messages.


Note: Select this option only when it is safe to rollback all the devices on the concerned storage array consistency group.

To understand the concept of forced rollback, refer to the following examples.



The above figure shows a simple application server where the application data (**/mnt1** and **/mnt2**) resides on five storage array devices (**D1**, **D2**, **D3**, **D4** and **D5**), which are a part of storage array consistency group. Consider that your NetBackup policy is configured such that you are using VSO snapshot method and backup selection leads to **/mnt1** and **/mnt2**. Once the snapshot is conducted, it will be conducted at consistency group level. So, now at the time of PIT Rollback Restore if you choose only **/mnt1** for such a restore, you will

get a warning stating partial selection. This is because of the fact that the snapshot being referred for restore has information about group snapshot consisting of devices D1, D2, D3, D4 and D5 where /mnt1 resides on D1, D2 and D3 only; and reverting back at consistency group will affect all devices in this group.



The above figure shows a simple application server where the application data (/mnt1 and /mnt2) resides on three devices (D1, D2, D3) which are part of a storage array consistency group. The consistency group is comprised of four devices (D1, D2, D3 and D4). Consider that your NetBackup policy is configured such that you are using VSO snapshot method and backup selection leads to /mnt1 and /mnt2. Once the snapshot is conducted, it will be conducted at consistency group level. So, now at the time of PIT Rollback Restore if you select both /mnt1 and /mnt2 for such a restore, the restore will fail due to mis-match of the devices. Because, the group consists of four devices whereas the devices being asked for restore are only three. Reverting back at consistency group will affect all devices in this group.

The rest of the procedure is identical to a normal restore as explained in the [NetBackup Backup, Archive, and Restore Getting Started Guide](#) and help.

Notes for restoring individual files from an Instant Recovery snapshot

The following items pertain to restoring individual files from an Instant Recovery snapshot:

- When you restore files from a snapshot that is made for an Instant Recovery off-host alternate client backup:
NetBackup consults the exclude list on the alternate client even when it restores files to the primary client. If the exclude list on the alternate client is different from the exclude list on the primary client, any files that are listed in the exclude list on the alternate client are not restored to the primary client.
For example, if the alternate client's exclude list has the entry *.jpg, and some .jpg files were included in the primary client backup, the .jpg files can be selected for the restore but are not in fact restored. To restore the files, you must change the exclude list on the alternate client.
- When you restore files from a snapshot that is made for an Instant Recovery backup (local or off-host alternate client):
If the exclude list is changed after the backup occurred, NetBackup honors the latest version of the exclude list during the restore. Any of the files that are listed in the current exclude list are not restored. Also, as noted in the previous item, the exclude list on the alternate client takes precedence over the exclude list on the primary client.
For example: If the current version of the exclude list has the entry *.jpg, and some .jpg files were included in the backup, the .jpg files can be selected for the restore but are not in fact restored. To restore the files, you must change the exclude list on the primary (or alternate) client.

Note: For ordinary backups (not based on snapshots), any files that were included in the exclude list are not backed up. For snapshot-based backups, however, all files are included in the snapshot. The exclude list is consulted only when a storage unit backup is created from the snapshot. If the snapshot is retained after the backup (for the Instant Recovery feature) and the snapshot is available at the time of the restore, NetBackup restores files from the snapshot. Since all files are available in the snapshot (including those that would be excluded from a storage unit backup), NetBackup incorrectly consults the current exclude list on the client or alternate client. Any files in the exclude list are skipped during the restore.

This issue will be addressed in a future release of NetBackup.


About configurations for restore

Snapshot Client backups can be restored in several ways, depending on your configuration.

About restoring over the LAN

Data can be restored from the storage device to the media server and from the media server over the LAN to the client. This kind of restore is also used for ordinary (non-Snapshot Client) backups.

Figure 11-1 Restore over LAN



The following table describes the phases that are illustrated in the diagram.


Phase	Action
Phase 1	Media server reads data from local storage.
Phase 2	Media server sends the data to the client over the LAN.
Phase 3	Client restores the data to disk (disk can be locally attached or on SAN).

About restoring over the SAN to a host acting as both client server and media server

This type of restore requires the `FORCE_RESTORE_MEDIA_SERVER` option in the server's `bp.conf` file.

See the [NetBackup Administrator's Guide, Volume I](#), for details on the `FORCE_RESTORE_MEDIA_SERVER` option.

Figure 11-2 Restore over the SAN



The following table describes the phases that are illustrated in the diagram.

Phase	Action
Phase 1	Client/media server reads data from tape over the SAN.
Phase 2	Client restores the data to disk (disk can be locally attached or on SAN). (Requires use of FORCE_RESTORE_MEDIA_SERVER option in bp.conf file.)

About restoring directly from a snapshot

If the **Keep snapshot after backup** option was turned on for the backup, the data can be restored from a mirror disk by restoring individual files from the snapshot, or by restoring the entire snapshot. This type of restore must be done from the command prompt. You can use a copy command such as UNIX `cp`.

Note the following:

- This type of restore cannot be done from the NetBackup Administration Console.
- This type of restore is not the Instant Recovery feature.

See [“About restoring from a disk snapshot”](#) on page 165.

You can restore individual files through the OpsCenter GUI. The restore is possible if the **Index From Snapshot** or **Backup From Snapshot** operation is selected while creating a storage lifecycle policy for snapshot replication.

About restoring from a disk snapshot

If the **Retain snapshot after backup** parameter is set to **Yes**, the snapshot is retained on the mirror disk after the backup completes. From the command line, you can restore individual files or the entire snapshot directly from the disk, rather than restoring from tape.

Note: Unless the backup was made with the Instant Recovery feature, you cannot restore from a snapshot by means of the Backup, Archive, and Restore interface. You must perform the restore manually at the command line.

About restoring on UNIX

The following procedures are for UNIX.

Restoring individual files on UNIX

To restore individual files, locate and mount the snapshot file system. Copy files from that file system using system commands, such as `cp` and `ftp`, as follows.

To restore individual files on UNIX

- 1 To list the identifiers of current snapshots, use the `bpfis` command with the `query` option:

```
/usr/opensv/netbackup/bin/bpfis query
```

This returns the ID (FIS IDs) of all current snapshots. For example:

```
INF - BACKUP START 3629
INF - FIS IDs: 1036458302
INF - EXIT STATUS 0: the requested operation was successfully
completed
```

- 2 For each snapshot identifier, enter `bpfis query` again, specifying the snapshot ID:

```
/usr/opensv/netbackup/bin/bpfis query -id 1036458302
```

This returns the path of the original file system (snapshot source) and the path of the snapshot file system. For example:

```
INF - BACKUP START 3629
INF - Snapshot image host : ricopico
INF - Snapshot image owner: GENERIC
INF - Time created       : Mon Oct  7 19:35:38 2002
INF - REMAP FILE BACKUP /mnt/ufscon USING

/tmp/_vrts_frzn_img_26808/mnt/ufscon

OPTIONS:ALT_PATH_PREFIX=/tmp/_vrts_frzn_img_26808,FITYPE=MIRROR,
MNTPOINT=/mnt/ufscon,FSTYPE=ufs
INF - EXIT STATUS 0: the requested operation was successfully
completed
```

In this example, the primary file system is `/mnt/ufscon` and the snapshot file system is `/tmp/_vrts_frzn_img_26808/mnt/ufscon`.

- 3 Copy the files from the mounted snapshot file system to the original file system.

Restoring the entire snapshot on UNIX

You can recover data from the disk snapshot in several ways, depending on your hardware configuration and the snapshot method that the policy used.

To restore the entire snapshot if the snapshot method was FlashSnap

- 1 Unmount the snapshot source (original file system) and the snapshot file system on the alternate client:

```
umount original_file_system
umount snapshot_image_file_system
```

To locate the file systems:

See [“Restoring individual files on UNIX”](#) on page 165.

- 2 Deport the snapshot on the alternate-client:

- Find the VxVM disk group:

```
vxvg list
```

The format of the disk group name is as follows:

```
SPLIT-primaryhost_diskgroup
```

If `vxvg list` does not show the disk group, the group might have been deported. You can discover all the disk groups, including deported ones, by entering:

```
vxdisk -o alldgs list
```

The disk groups in parentheses are not imported on the local system.

- Deport the VxVM disk group:

```
vxvg deport SPLIT-primaryhost_diskgroup
```

- 3 Import and join the VxVM disk group on the primary (original) client:

```
vxvg import SPLIT-primaryhost_diskgroup
vxrecover -g SPLIT-primaryhost_diskgroup -m
vxvg join SPLIT-primaryhost_diskgroup diskgroup
```

- 4 Start the volume and snap back the snapshot volume as follows, using the `-o resyncfromreplica` option:

```
vxvol -g SPLIT-primaryhost_diskgroup start SNAP_diskgroup_volume
vxassist -g SPLIT-primaryhost_diskgroup -o resyncfromreplica
snapback SNAP_diskgroup_volume
```

To restore the entire secondary disk if the snapshot was made on an EMC, Hitachi, or HP disk array

WITH CAUTION, you can use hardware-level restore to restore the entire mirror or secondary disk to the primary disk.

If the disk is shared by more than one file system or VxVM volume, there may be unintended results. Read the following:

Warning: Hardware-level disk-restore (such as with the `symmir -restore` command) can cause data loss if more than one file system or more than one VxVM volume shares the primary disk. The hardware-level restore overwrites the entire primary disk with the contents of the mirror disk.

This overwrite can be a problem if you attempt to restore a snapshot of one of the file systems or one of the VxVM volumes that share the same disk. The other file systems or volumes sharing the disk may have older data that you do not want to write back to the primary. When the hardware-level disk-restore takes place, the older data replaces the newer data on the primary disk.

About restoring on Windows

The following procedures are for Windows.

Restoring individual files on Windows

Use the following procedure to restore individual files on Windows.

To restore individual files on Windows

- 1 To list the identifiers of current snapshots, use the `bpfis` command with the `query` option:

```
/usr/opensv/netbackup/bin/bpfis query
```

This returns the ID (FIS IDs) of all current snapshots. For example:

```
INF - BACKUP START 3629
INF - FIS IDs: 1036458302
INF - EXIT STATUS 0: the requested operation was successfully
completed
```

- 2 For each snapshot identifier, enter `bpfis query` again, specifying the snapshot ID:

```
/usr/opensv/netbackup/bin/bpfis query -id 1036458302
```

This returns the path or the original file system (snapshot source) and the GUID (Global Universal Identifier) representing the snapshot volume. For example:

```
INF - BACKUP START 2708
INF - Snapshot method: FlashSnap
INF - Snapshot image host : tire
INF - Snapshot image owner : NBU
INF - Time created : Sat Oct 25 15:26:04 2003

INF - REMAP FILE BACKUP H:\ USING
\\?\Volume{54aa666f-0547-11d8-b023-00065bde58d1}\
OPTIONS:ALT_PATH_PREFIX=C:\Program Files\Veritas\NetBackup\
Temp\_vrts_frzn_img_2408,FITYPE=MIRROR,MNTPPOINT=H:\,FSTYPE=NTFS
INF - EXIT STATUS 0: the requested operation was successfully
completed
```

In this example the snapshot file system is `H:\` and the GUID is `\\?\Volume{54aa666f-0547-11d8-b023-00065bde58d1}\`.

- 3 To restore individual files from the snapshot volume:

- Mount the GUID to an empty NTFS directory:

```
mountvol C:\Temp\Mount
\\?\Volume{54aa666f-0547-11d8-b023-00065bde58d1}\
```

- Copy the file to be restored from the temporary snapshot mount point (in this example, `C:\Temp\Mount`) to the primary volume.

Restoring the entire snapshot on Windows

The following procedure is applicable if the snapshot method was FlashSnap.

To restore the entire snapshot on Windows

1 Deport the snapshot on the alternate-client:

- Find the VxVM disk group:

```
vxdbg list
```

The format of the disk group name is as follows:

```
SPLIT-primaryhost_diskgroup
```

- Deport the VxVM disk group:

```
vxdbg -g split_diskgroup deport
```

2 Import and join the VxVM disk group on the primary (original) client:

```
vxassist rescan  
vxdbg -g split_diskgroup import  
vxdbg -g split_diskgroup -n diskgroup join
```

3 Snap back the snapshot volume, using the `-o resyncfromreplica` option:

```
vxassist -o resyncfromreplica snapback \Device\HarddiskDmVolumes\diskgroup\snap_volume
```

Snapshot management

This chapter includes the following topics:

- [About NetBackup and CloudPoint integration for snapshot management](#)
- [Troubleshooting Snapshot Client and NetBackup integration issues](#)

About NetBackup and CloudPoint integration for snapshot management

NetBackup interfaces with snapshot management servers to manage arrays for discovery of devices and snapshot lifecycle management. Snapshot management servers are installed separately apart from NetBackup and they have plugins for various storage arrays. Veritas CloudPoint (TM) is an example of a snapshot management server for NetBackup. Using the NetBackup Snapshot management framework, you can use CloudPoint for taking snapshots of your images. You can protect all the on-premises storage arrays that are supported by CloudPoint.

The following table describes the underlying tasks.

Table 12-1 Tasks for snapshot management

Task	Description
Configure the CloudPoint server in NetBackup.	You can configure the CloudPoint server as a snapshot management server. To configure the CloudPoint server in NetBackup you need to add the credentials of the CloudPoint server.

Table 12-1 Tasks for snapshot management (*continued*)

Task	Description
Configure the CloudPoint plug-ins in NetBackup.	The CloudPoint plug-ins that are installed on the CloudPoint server must be registered and configured in NetBackup with the associated CloudPoint server.
Configure a Standard policy to use VSO snapshot method.	See “Configuring a Snapshot Client policy” on page 39. See “Selecting the snapshot method” on page 43.
Backup and restore procedures	See “About performing a backup” on page 155. See “About performing a restore” on page 156.

Considerations

Consider the following when integrating NetBackup with CloudPoint:

- Deletion of CloudPoint host entry and its associated plug-ins is not supported in NetBackup.
Using CloudPoint, if you delete plug-ins that are configured in NetBackup, the images of snapshots that are taken using CloudPoint will be unusable.
- Back up from snapshot and index from snapshot is supported. Replication operations are not supported.
- For snapshot and index for snapshot, you can use any available storage unit.
- Post integration, all the related operations must be performed from NetBackup. Results of operations that are performed outside NetBackup, will not be visible in NetBackup. CloudPoint has its own RBAC, where NetBackup is one of the users and operations performed only through NetBackup will be visible in NetBackup. For example, even though you can add CloudPoint plug-in from CloudPoint, you must add the plug-in from NetBackup, or the plug-in may not be visible in NetBackup.
- HP-UX Native volume group Version 2.0 and later is not supported.
- Raw device and raw partitions are not supported.
- File system verification is not supported. You must skip the verification step and perform a force rollback. In the **Restore Marked Files** dialog box, select the **Skip verification and force rollback** check box.
- Consider all the CloudPoint limitations.
Refer to the *Veritas CloudPoint Installation and Upgrade Guide*.

- NetBackup integration is not supported with the CloudPoint freemium version.
- It is recommended not to reinstall the CloudPoint server but to upgrade it. However, if you reinstall the CloudPoint server you need to reconfigure the CloudPoint server and the associated CloudPoint plug-ins.

- Perform the following tasks for RHEL:
In the recent RHEL version, there are changes in usage of partition delimiter. Based on last character in the device path, OS decides if partition delimiter must be added or not. To address the change, update the udev rule to add partition delimiter 'p' on RHEL

- Open the `/lib/udev/rules.d/62-multipath.rules` file.

- Update the existing line:

```
RUN+=" $env{MPATH_SBIN_PATH}/kpartx -an $temprnode", \
ENV{DM_MULTIPATH_NEED_KPARTX}=""
```

To

```
RUN+=" $env{MPATH_SBIN_PATH}/kpartx -an -p p $temprnode", \
ENV{DM_MULTIPATH_NEED_KPARTX}=""
```

- Restart the multipath service. Run the following command:

```
# service multipathd restart
```

- For devices in LVM, if multiple paths are used, ensure that **multipathing** service is enabled.
- On SUSE Linux, ensure that the **automount** service is disabled.
- Junction mount point on filler or array is not supported for snapshot operations with VSO FIM snapshot method.
- For client-side nested mount points, rollback of parent mount points is not supported.

Troubleshooting Snapshot Client and NetBackup integration issues

For troubleshooting also consider the limitations.

Policy validation fails with error 130

Policy validation fails with error 130, due to one of following reasons:

- Certificate mismatch on client.

Workaround

Remove the CloudPoint CA certificate from the following location and validate the policy again.

`install_location/var/global/cloudpoint`

- Cloud point server is not accessible.

Policy validation fails with error 223

This error is encountered if the following options are not configured:

- In the **Snapshot Client and Replication Director** dialog box, select the **Retain snapshot for Instant Recovery or SLP management** option.
- In the **Schedules** tab, select **Snapshots only** option.

Policy validation fails with error 4201

Policy validation fails with error 4201, due to one of following reasons:

- The source disk is not associated with configured CloudPoint server.

Workaround

Check if you have configured the CloudPoint plug-in associated with a particular disk asset.

- CloudPoint server takes additional time to discover.

Workaround

Wait for sometime so that all the devices from array are discovered.

- CloudPoint server is not accessible from Master or Client.

Or

Certificate mismatch on client.

Or

CloudPoint certificate has been expired.

Workaround

Refresh the CloudPoint certificates on client.

Certificate location: `/usr/opensv/var/global/cloudpoint/<CloudPoint server name>.pem`

- HP3 par is loaded and there are multiple WS API operations are running.

Workaround

Try snapshot creation after sometime.

VSO policy validation fails with error Asset no_vso_device not found

VSO policy validation fails with the following error:

```
CVSOPugin::selectSnapshots : getAssetDetails failed with status:
```

```
[error], Error:[Asset no_vso_device not found]
```

Workaround

When you encounter this error, run the `lsscsi` command to verify if the device selected in the backup selection is available on the client.

If the device is not listed, verify the connectivity issue between the array and host or get in touch with your system administrator.

Backup fails with PFI rotation error 13

Backup fails with the following error:

PFI rotation error: 13

The errors occurs because the operation takes longer to complete due to network connectivity, causing jobs to time out and fail.

Workaround

Increase the `REQUEST_DELIVERY_TIMEOUT` configuration option from the default of 300 (5 minutes) to more than 3600 seconds (more than 60 minutes).

This option does not appear in the NetBackup Administration Console host properties. See the *NetBackup Commands Reference Guide* for information about using the `bpgetconfig` and the `bpsetconfig` commands to change the configuration option in the `bp.conf` file (UNIX) or the registry (Windows).

For HP 3PAR devices, snapshot restore fails with error 2800

If there is any clone snapshot or physical copy associated with source device, snapshot restore fails.

Workaround

Break the clone relationship or delete the clone snapshot.

Troubleshooting

This chapter includes the following topics:

- [About gathering information and checking logs](#)
- [Logging directories for UNIX platforms](#)
- [Logging folders for Windows platforms](#)
- [Customer support contact information](#)
- [Latest patches and updates](#)
- [Snapshot provider information](#)
- [Important notes on &newProductName;](#)
- [&newProductName; installation problems](#)
- [FlashBackup and status code 13](#)
- [Single file restore from a FlashBackup Instant Recovery snapshot of a file protected by Windows VSS writer fails](#)
- [Identifying and removing a left-over snapshot](#)
- [Removing a VxVM volume clone](#)
- [Alternate client restore and backup from a snapshot fails](#)
- [Restore from a snapshot fails with status 2800](#)
- [Raw Partition restore fails with the message 'FlashBackup-Windows policy restore error'](#)
- [Snapshot creation fails with error 156](#)
- [Snapshot fails with error 20](#)

- Restore from a snapshot fails with status 133
- Backup from snapshot fails with error 50
- Backup from snapshot parent job fails with error 4213: Snapshot import failed
- Snapshot job fails and the snapshot command does not recognize the volume name
- Snapshot creation fails with error 4220
- Snapshot creation fails when the same volume is mounted on multiple mount points of the same host
- Snapshot-based backup and restore failure
- Multiple snapshot jobs fail with code 156 or 1541.
- FlashBackup policy fails, with multiple backup selections [Cache =]
- Partial backup failure with 'Snapshot encountered error 156'
- Backup of file system validation fails with error 223
- Policy validation fails if the specified CIFS share path contains a forward slash
- Policy validation fails with status code 223
- An NDMP snapshot policy for wildcard backup fails with error 4201
- Troubleshooting with bpfis log
- Limitations of using HP-UX 11.31
- While configuring a policy for snapshot operations, the NetBackup Administration Console hangs
- Accelerator enabled incremental backup of NetApp NAS volume
- Snapshot method: Auto

About gathering information and checking logs

You can resolve many problems on your own by creating logging directories, reproducing the problem, and checking the logs.

For an in-depth description of NetBackup logs, refer to the [NetBackup Troubleshooting Guide](#).

For explanations of NetBackup job status codes, refer to the [NetBackup Status codes Reference Guide](#).

Logging directories for UNIX platforms

Snapshot Client backup messages are written to the directories that are shown in See [Table 13-1](#) on page 178., if the directories exist.

Note the following:

- To create detailed log information, place a VERBOSE entry in the `bp.conf` file on the NetBackup master and client. Or set the Global logging level to a high value in the **Logging** dialog, under both **Master Server Properties** and **Client Properties**.
- These directories can eventually require a lot of disk space. Delete them when you are finished troubleshooting and remove the VERBOSE option from the `bp.conf` file. Or reset the Global logging level to a lower value.
- Messages pertaining to NAS_Snapshot can be found in the ndmp unified log (originator ID 151), in `/usr/opensv/logs`.
- The bpfis log directory contains the bpfis and ostfi (OST plug-in) logs.

UNIX logging directories for backup

During a backup, &newProductName; messages are logged to the following directories. Create these directories using an access mode of 755 so NetBackup can write to the logs.

You can use the `/usr/opensv/netbackup/logs/mklogdir` script to create these directories.

Table 13-1 UNIX logging directories for backup

Path of log directory	Where to create the directory
<code>/usr/opensv/netbackup/logs/bpbrm</code>	NetBackup master server if Instant Recovery backup is set to snapshot only; otherwise, on media server
<code>/usr/opensv/netbackup/logs/bptm</code>	NetBackup media server
<code>/usr/opensv/netbackup/logs/bpbkar</code>	NetBackup client or alternate client
<code>/usr/opensv/netbackup/logs/bpfis</code>	NetBackup client or alternate client
<code>/usr/opensv/netbackup/logs/bppfi</code>	NetBackup client or alternate client

UNIX logging directories for restore

During a restore, &newProductName; messages are logged to the following directories on the master server. Create these directories using an access mode of 755.

Table 13-2 UNIX logging directories for restore

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bprestore	NetBackup master server
/usr/opensv/netbackup/logs/bprd	NetBackup master server
/usr/opensv/netbackup/logs/bpbmm	NetBackup master server if Instant Recovery backup is set to snapshot only; otherwise, on media server
/usr/opensv/netbackup/logs/bptm	NetBackup media server

snapctl driver messages

Messages from the `snapctl` driver are logged in the client's `/var/adm/messages` file along with other kernel messages.

Logging folders for Windows platforms

During a backup, &newProductName; messages are written to the folders that are listed in See [Table 13-3](#) on page 180. if the folders exist. You can use the following command to create these folders:

```
install_path\NetBackup\logs\mklogdir.bat
```

The default path for the logs is the following:

```
C:\Program Files\Veritas\NetBackup\logs
```

Since a different path can be set during installation, the paths that are listed in this topic are `install_path\NetBackup\logs`.

Note: To create detailed logs, set the Global logging level to a high value, in the **Logging** dialog, under both **Master Server Properties** and **Client Properties**.

The log folders can eventually require a lot of disk space. Delete them when you are finished troubleshooting and set the logging level on master and client to a lower value.

Windows logging folders for backup

During a backup, &newProductName; messages are logged to the following folders.

Messages pertaining to NAS_Snapshot can be found in the ndmp unified log (originator ID 151), in *install_path\NetBackup\logs*.

Table 13-3 Windows logging folders for backup

Path of log directory	Where folder is created
<i>install_path\NetBackup\logs\bpbrm</i>	NetBackup master server if Instant Recovery backup is set to snapshot only; otherwise, on media server
<i>install_path\NetBackup\logs\bptm</i>	NetBackup media server
<i>install_path\NetBackup\logs\bppfi</i>	NetBackup client or alternate client
<i>install_path\NetBackup\logs\bpbkar</i>	NetBackup client or alternate client
<i>install_path\NetBackup\logs\bpfis</i>	NetBackup client or alternate client

Windows logging folders for restore

During a restore, &newProductName; messages are logged to the following folders on the master server.

Table 13-4 Windows logging folders for restore

Path of log directory	Where folder is created
<i>install_path\NetBackup\logs\bprestore</i>	NetBackup master server
<i>install_path\NetBackup\logs\bprrd</i>	NetBackup master server
<i>install_path\NetBackup\logs\bpbrm</i>	NetBackup master server if Instant Recovery backup is set to snapshot only; otherwise, on media server
<i>install_path\NetBackup\logs\bptm</i>	NetBackup media server

Configuring VxMS logging

The following procedures describe how to configure VxMS logging for NetBackup.

Except as noted in this topic, you can also use the Logging Assistant (in the NetBackup Administration Console) to configure VxMS logging.

VxMS logging on a NetBackup UNIX or Linux client

To configure VxMS logging on a NetBackup UNIX or Linux client

- 1 Create the VxMS log directory:

```
/usr/opensv/netbackup/logs/vxms
```

Note: For logging to occur, the VxMS directory must exist.

Note: If you have run the NetBackup `mklogdir` command, the VxMS log directory already exists.

- 2 Add the following to the `/usr/opensv/netbackup/bp.conf` file:

```
VXMS_VERBOSE=<numeric value of 0 or greater>
```

See [Table 13-5](#) for the available logging levels.

- 3 To change the log location, enter the following in the `bp.conf` file:

```
vxmslogdir=path to new log location
```

Note: If the VxMS log location is changed, the Logging Assistant does not collect the logs.

VxMS logging on a NetBackup Windows client

To configure VxMS logging on a NetBackup Windows client

- 1 Create the VxMS log directory:

```
install_path\NetBackup\logs\vxms
```

Note: For logging to occur, the VxMS folder must exist.

Note: If you have run the NetBackup `mklogdir.bat` command, the VxMS log directory already exists.

- 2 In the Windows registry, create the DWORD registry entry `VXMS_VERBOSE` in the following location:

HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion > Config

- 3 To configure the logging level, set the numeric value of `VXMS_VERBOSE` to 0 or greater. Larger numbers result in more verbose logs.

See [Table 13-5](#) for the available logging levels.

- 4 To change the log location:

- Open regedit and go to the following location:
HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion
- Create the registry entry `vxmslogdir` with a string value (`REG_SZ`). For the string value, specify the full path to an existing folder.

Note: You can use NTFS compression on VxMS log folders to compress the log size. The new logs are written in compressed form only.

Note: If the VxMS log location is changed, the Logging Assistant does not collect the logs.

VxMS logging levels

[Table 13-5](#) lists the VxMS logging levels.

Note: Logging levels higher than 5 cannot be set in the Logging Assistant.

Note: Logging levels higher than 5 should be used in very unusual cases only. At that level, the log files and metadata dumps may place significant demands on disk space and host performance.

Table 13-5 VxMS logging levels

Level	Description
0	No logging.
1	Error logging.
2	Level 1 + warning messages.
3	Level 2 + informative messages.
4	Same as level 3.
5	Highly verbose (includes level 1) + auxiliary evidence files (.mmf, .dump, VDDK logs, .xml, .rvpmem). You can set the logging level for the VDDK messages.
6	VIX (VMware virtual machine metadata) dumps only.
7	VHD (Hyper-V virtual machine metadata) dumps only.
> 7	Full verbose + level 5 + level 6 + level 7.

Customer support contact information

Before calling customer support, gather as much log information as possible. Be sure to have the following information ready:

- NetBackup version
- Operating system version of the NetBackup master and media server and NetBackup &newProductName; client
- Note whether or not the action that failed had ever worked and whether the problem is repeatable
- Log information

Latest patches and updates

For other Veritas products such as the Veritas File System and Volume Manager, or Storage Foundation, install the latest patches and updates for those products. Installing the latest software can fix a variety of issues.

For example:

If you receive status code 156 on a FlashSnap alternate client backup, and the client data is configured in Storage Foundations for Windows 4.1, note: the volume to be backed up may not appear on the alternate client. This situation can occur even though the disk group has been split on the primary client and the disk was deported to the alternate client. Installing the latest Storage Foundation patches has been known to correct this problem.

Snapshot provider information

Although NetBackup communicates with the snapshot provider on its own (you do not configure providers), the NetBackup logs occasionally refer to providers by name. A knowledge of which provider was used can help in understanding the snapshot process and in correcting problems.

Important notes on &newProductName;

Note the following:

- If backup or restore jobs are running slowly, verify that the network interface cards (NIC) are set to full duplex. Half duplex often causes poor performance. For assistance to view and reset duplex mode for a particular host or device, consult the manufacturer's documentation. You may be able to use the `ifconfig` (or `ipconfig`) command to view and reset duplex mode, as explained in the *NetBackup Troubleshooting Guide*.
- Inconsistency in the count of estimated files in the snapshot job and the restore job details. In the snapshot job details the estimated files field displays a ' ' (blank) whereas in the restore job details the estimated files field displays the value as '1'.
- The disk containing the client's data must be a SCSI or Fibre Channel device if you use NetBackup Media Server or Third-Party Copy Device.
- The disk containing the client's data must be visible to both the client and the media server, for NetBackup Media Server or Third-Party Copy Device. The disk can be connected through SCSI or Fibre Channel.

- For NetBackup Media Server or Third-Party Copy Device, a disk must support serialization or support SCSI Inquiry Page Code 83.
- For Third-Party Copy Device or NetBackup Media Server, a particular storage unit or group of storage units must be specified for the policy. Do not choose **Any_available**. Configuration instructions are available: See “[Configuring a Snapshot Client policy](#)” on page 39.
- The *storage_unit_name* portion of a `mover.conf.storage_unit_name` file name must exactly match the storage unit name (such as `nut-4mm-robot-tl14-0`) that you have defined for the policy.
Help is available for creating a `mover.conf.storage_unit_name` file. See the *NetBackup Snapshot Client Configuration* document:
<http://www.veritas.com/docs/000081320>
Similarly, the *policy_name* portion of a `mover.conf.policy_name` file name must match the name of the policy that the third-party copy device is associated with.
- For the legacy disk array methods (TimeFinder, ShadowImage, or BusinessCopy), the client data must reside in a device group. The data must be on the primary disk and be synchronized with a mirror disk. Assistance from the disk array vendor may also be required.
For information on disk configuration requirements for the legacy array methods, see the *NetBackup Snapshot Client Configuration* document:
<http://www.veritas.com/docs/000081320>
- If the **Keep snapshot after backup** option is changed from **yes** to **no**, the last snapshot that is created for that policy must be deleted manually before the backup is rerun. Use the `bpfis` command to delete the snapshot. Refer to the man page for `bpfis`.
- During a third-party copy device backup, if tape performance is slow, increase the buffer size. To do so, create one of the following files on the media server:

```
/usr/openv/netbackup/db/config/SIZE_DATA_BUFFERS_TPC.policy_name
/usr/openv/netbackup/db/config/SIZE_DATA_BUFFERS_TPC.storage_unit_name
/usr/openv/netbackup/db/config/SIZE_DATA_BUFFERS_TPC
```

For third-party copy backup, the size of the data buffer is 65536 bytes (64K), by default. To increase it, put a larger integer in the `SIZE_DATA_BUFFERS_TPC` file. For a buffer size of 96K, put 98304 in the file. If not an exact multiple of 1024, the value that is read from the file is rounded up to a multiple of 1024. The file name with no extension (`SIZE_DATA_BUFFERS_TPC`) applies as a default to all third-party copy backups, if neither of the other file-name types exists. A `SIZE_DATA_BUFFERS_TPC` file with the *.policy_name* extension applies to backups that the named policy runs. The *.storage_unit_name* extension

applies to backups that use the named storage unit. If more than one of these files applies to a given backup, the buffers value is selected in this order:

SIZE_DATA_BUFFERS_TPC.*policy_name*
 SIZE_DATA_BUFFERS_TPC.*storage_unit_name*
 SIZE_DATA_BUFFERS_TPC

As soon as one of these files is located, its value is used. A *.policy_name* file that matches the name of the executed policy overrides the value in both the *.storage_unit_name* file and the file with no extension. The *.storage_unit_name* file overrides the value in the file with no extension.

You can set the maximum buffer size that a particular third-party copy device can support.

A third-party copy device is not used if it cannot handle the buffer size that is set for the backup.

- Replication Director is a feature that makes use of an OpenStorage application to manage snapshot replication. The snapshots are stored on the storage systems of partnering companies.

Note: Replication Director uses the NetApp DataFabric Manager server for data movement and not the media server as in most cases.

&newProductName; installation problems

If you receive the following message during installation:

```
/usr/opensv/netbackup/bin/version not found.  
Add-On Product Installation Aborted.
```

you have tried to install the &newProductName; software before you install the base NetBackup software.

FlashBackup and status code 13

Status 13 can result from any of the following:

The FlashBackup cache partition may have run out of space

In that case, the cache partition may not be large enough for the backup requirements.

If the cache partition is full, messages such as the following appear in the system log:

Single file restore from a FlashBackup Instant Recovery snapshot of a file protected by Windows VSS writer fails

```
WARNING: sn_alloccache: cache /dev/vx/rdisk/flashldg/f full - all
snaps using this cache are now unusable
WARNING: sn_failsnap: snapshot id 5 failed error 28
```

Specify a larger cache partition, or designate additional cache partitions in the **Backup Selections** list.

See [“Requirements for the cache partition”](#) on page 80.

Removing stale snapshots (Solaris)

On Solaris, if your cache partition runs out of space, stale snapshots may be taking up space on the cache partition. Stale snapshots are those that FlashBackup did not automatically delete.

To remove stale snapshots (Solaris)

- 1 Determine if there are stale snapshots on your Solaris client by executing the following:

```
/usr/opensv/netbackup/bin/driver/snaplist
```

- 2 For each snapshot that is listed, run the following to make sure a `bbpkar` process is associated with it:

```
ps -eaf |grep ident
```

where *ident* is the snapshot process id displayed by the `snaplist` command.

- 3 Remove snapshots that do not have an associated `bbpkar` process by entering the following:

```
/usr/opensv/netbackup/bin/driver/snapoff snapn
```

where *snapn* is the snapshot id displayed by the `snaplist` command.

Single file restore from a FlashBackup Instant Recovery snapshot of a file protected by Windows VSS writer fails

When performing a single file recovery from a FlashBackup Instant Recovery snapshot of a file that is protected by Windows VSS writer, the restore completes successfully however the file is not restored.

When a file is protected by a VSS writer, it should be backed up and restored with VSS writer involvement. If the VSS snapshot that is taken for FlashBackup does

not have VSS writer involvement, the data in the file may not be consistent. This file should not be used for restore from either tape or snapshot.

A proper way to protect this file is to have a file backup with Shadow Copy Component file list directive specified. In this case, a VSS snapshot with VSS writer is taken.

Identifying and removing a left-over snapshot

NetBackup ordinarily removes snapshots after the &newProductName; backup completes, unless the **Keep snapshot after backup** parameter was set to **Yes**. However, as a result of some system failures, such as a system crash or abnormal backup termination, the snapshot may not be removed.

To identify and remove a left-over snapshot

- 1 Use the `bpfis` command with the `query` option to list the current snapshots. Enter the following on the client or alternate client, depending on the type of backup:

```
/usr/opensv/netbackup/bin/bpfis query
```

This command returns the IDs (FIS IDs) of all current snapshots. For example:

```
INF - BACKUP START 3629
INF - FIS IDs: 1036458302
INF - EXIT STATUS 0: the requested operation was successfully
completed
```

In this example, the snapshot ID is 1036458302.

- 2 If the `bpfis` output shows the ID of the snapshot, delete it as follows:

```
bpfis delete -id snapshot_id
```

If `bpfis` removed the snapshot, you can skip the rest of this procedure.

- 3 Solaris, HP, AIX, Linux: if bpfis could not remove the snapshot, enter the following (on the client or alternate client) when no backups are running:

```
df -k
```

This command displays all mounted file systems, including any snapshots of a mounted file system.

If a snapshot backup is currently running, the snapshot should not be deleted. NetBackup deletes it when the backup completes.

Here are two snapshots from a `df -k` listing:

```
/dev/dsk/clt3d2s4 1048800 73076 914742 8% /tmp/_vrts_frzn_img_wil_vxfs_1299000
/dev/vx/dsk/clone_qes_clone/ufs 38383 21678 12867 63% /tmp/_vrts_frzn_img
__mix_ufs_1299000
```

The snapshot appears in the following form:

```
/tmp/_vrts_frzn_img__filesystemname_pid
```

- 4 Solaris, HP, AIX, Linux: unmount the unneeded snapshot file systems (on the client or alternate client, depending on the type of backup).

The next step depends on the type of snapshot.

- 5 For `nbu_snap` (Solaris only):

- Enter the following to display leftover snaps:

```
/usr/opensv/netbackup/bin/driver/snaplist
```

- To remove a leftover snap, enter

```
/usr/opensv/netbackup/bin/driver/snapoff snap1 ... snapn
```

More information is available on `snaplist` and `snapoff`.

See [“About managing nbu_snap”](#) on page 206.

- 6 For VxVM (Solaris, HP, AIX, Linux) and VVR (Solaris and HP):

Do the following on the client for VxVM, and on the alternate client for VVR:

- Enter the following to display unsynchronized mirror disks:

```
vxprint -g diskgroup
```

- Enter the following to resynchronize the mirror disks:

```
vxassist -g diskgroup -v volume snapback
```

7 For VxVM (Windows):

- Enter the following to display unsynchronized mirror disks:

```
vxdbg -g diskgroup dginfo
```

- Enter the following to resynchronize the mirror disks:

```
vxassist snapback \Device\HarddiskDmVolumes\diskgroup\snap_volume
```

8 For VxFS_Checkpoint (Solaris, HP, AIX, Linux):

- Enter the following VxFS command to display the name of the checkpoint:

```
/usr/lib/fs/vxfs/fscckptadm list /file_system
```

Note: *file_system* is the mount point of the primary file system that was backed up, NOT the snapshot file system that was unmounted in a previous step.

For example, if the snapshot file system that was unmounted is the following:

```
/tmp/_vrts_frzn_img__vm2_1765
```

the original file system, which should be specified on the `fscckptadm list` command, is the following:

```
/vm2
```

Example entry:

```
/usr/lib/fs/vxfs/fscckptadm list /vm2
```

Output:

```
/vm2
NBU+2004.04.02.10h53m22s:
    ctime           =  Fri Apr 02 10:53:23 2004
    mtime           =  Fri Apr 02 10:53:23 2004
    flags           =  removable
```

In this example, the name of the checkpoint is NBU+2004.04.02.10h53m22s.

- Remove the checkpoint by entering the following:

```
/usr/lib/fs/vxfs/fscckptadm remove name_of_checkpoint /file_system
```

For example:

```
/usr/lib/fs/vxfs/fsckptadm remove NBU+2004.04.02.10h53m22s /vm2
```

- If the checkpoint cannot be removed, unmount it (`umount`) and retry the following:

```
/usr/lib/fs/vxfs/fsckptadm remove name_of_checkpoint /file_system
```

- For more detail on removing VxFS clones, refer to the recommended actions for NetBackup status code 156 in the [NetBackup Troubleshooting Guide](#).

9 For TimeFinder, ShadowImage, BusinessCopy (Solaris or HP only):

Do the following on the client or alternate client, depending on the type of backup:

- To discover and remove any VxVM clones:
See [“Removing a VxVM volume clone”](#) on page 193.
- Enter the following to resynchronize the mirror disks:
For EMC arrays (TimeFinder):

```
symmir -g device_group establish LdevName
```

where *LdevName* is the logical device name of the standard device. For Hitachi and HP arrays (ShadowImage, BusinessCopy):

```
pairresync -g groupname -d dev_name
```

For more information about EMC, Hitachi, and HP arrays and resynchronizing disks, see the *NetBackup Snapshot Client Configuration* document:

<http://www.veritas.com/docs/000081320>

10 For VxFS_Snapshot (Solaris or HP only):

Using the mounted file system from a previous step, unmount the snapshot as follows:

```
umount -F vxfs /tmp/_vrts_frzn_img__filesystemname_pid
```

11 For FlashSnap (Solaris, HP, AIX, Linux):

Do the following on the client or alternate client, depending on the type of backup:

- Find the VxVM disk group:

```
vxvg list
```

- The format of the disk group name is as follows:

SPLIT-primaryhost_diskgroup

If `vxvg list` does not show the disk group, the group might have been deported. You can discover all the disk groups, including deported ones, by entering:

```
vxdisk -o alldgs list
```

The disk groups in parentheses are not imported on the local system.

- Deport the VxVM disk group:

```
vxvg deport SPLIT-primaryhost_diskgroup
```

- On the primary (original) client, import and join the VxVM disk group:

```
vxvg import SPLIT-primaryhost_diskgroup
vxrecover -g SPLIT-primaryhost_diskgroup -m
vxvg join SPLIT-primaryhost_diskgroup diskgroup
```

- On the primary (original) client, start the volume and snap back the snapshot volume:

```
vxvol -g SPLIT-primaryhost_diskgroup start SNAP-diskgroup_volume
vxassist snapback SNAP-diskgroup_volume
```

Example:

In this example, `chime` is the primary client and `rico` is the alternate client.

`lhddg` is the name of the original disk group on `chime`.

`chime_lhddg` is the split group that was imported on `rico` and must be rejoined to the original group on the primary `chime`.

On alternate client `rico`, enter:

```
vxvg deport chime_lhddg
```

On primary client `chime`, enter:

```
vxvg import chime_lhddg
vxrecover -g chime_lhddg -m
vxvg join chime_lhddg lhddg
vxvol start SNAP-lhddg-vol01
vxassist snapback SNAP-lhddg-vol01
```

12 For FlashSnap (Windows):

- Find the VxVM disk group:

```
vxdbg list
```

- The format of the disk group name is as follows:

```
SPLIT-primaryhost_diskgroup
```

- Deport the VxVM disk group:

```
vxdbg -g split_diskgroup deport
```

- On the primary (original) client, import and join the VxVM disk group:

```
vxassist rescan
vxdbg -g split_diskgroup import
vxdbg -g split_diskgroup -n diskgroup join
```

- On the primary (original) client, snap back the snapshot volume:

```
vxassist snapback \Device\HarddiskDmVolumes\diskgroup\snap_volume
```

Removing a VxVM volume clone

A VxVM volume clone is a form of snapshot that might need manual deletion. For a description of disk clones, see the *NetBackup Snapshot Client Configuration* document:

<http://www.veritas.com/docs/000081320>

Major system interruptions, such as a system crash or unexpected restart, can prevent NetBackup from removing the clone. If the clone is not removed, subsequent backups of the client's data fail. Examine the `/usr/openv/netbackup/logs/bpfis` log for text such as the following:

```
19:13:07.686 [14981] <2> onlfi_vfms_logf: INF - do_cmd:
Command failed with status=20:
/usr/openv/netbackup/bin/bpdgclone -g wil_test -n vol01 -f /var/tmp/HDSTFCAAs7aOqD
</dev/null >/var/tmp/VfMSAAQ7aOqD 2>/var/tmp/VfMSBAAr7aOqD
19:13:07.687 [14981] <2> onlfi_vfms_logf: INF -
--- Dumping file /var/tmp/VfMSAAQ7aOqD (stdout):
19:13:07.687 [14981] <2> onlfi_vfms_logf: INF -
--- End of file /var/tmp/VfMSAAQ7aOqD
19:13:07.687 [14981] <2> onlfi_vfms_logf: INF -
```

```
--- Dumping file /var/tmp/VfMSBAAr7aOqD (stderr):
19:13:07.687 [14981] <2> onlfi_vfms_logf: INF - clone group and volume already exists
19:13:07.688 [14981] <2> onlfi_vfms_logf: INF - --- End of file /var/tmp/VfMSBAAr7aOqD
```

In this case, you must use the `bpdgclone` command with the `-c` option to remove the clone. Then resynchronize the mirror disk with the primary disk.

The following commands should be run on the client or alternate client, depending on the type of backup.

To remove a VxVM volume clone

- 1 When no backups are running, use the following VxVM command to list any clones.

```
vxvg list
```

If a backup configured with an array-specific snapshot method is currently running, a clone for that backup appears in the `vxvg` output. Do not delete the clone; NetBackup deletes it when the backup completes.

Example `vxvg` output:

NAME	STATE	ID
rootdg	enabled	983299491.1025.turnip
VolMgr	enabled	995995264.8366.turnip
wil_test_clone	enabled	1010532924.21462.turnip
wil_test	enabled	983815798.1417.turnip

In this example, the name suffix indicates `wil_test_clone` was created for a snapshot backup that was configured with an array-specific snapshot method. If a backup failed with log entries similar to those in this example, the clone must be manually deleted.

- 2 To remove the clone, enter the following:

```
/usr/opensv/netbackup/bin/bpdgclone -g disk_group -n volume -c clone
```

For the previous example, you would enter:

```
/usr/opensv/netbackup/bin/bpdgclone -g wil_test -n vol01 -c wil_test_clone
```

where `wil_test` is the name of the disk group, `vol01` is the name of the VxVM volume, and `wil_test_clone` is the name of the clone. Use the Volume Manager `vxprint` command to display volume names and other volume information.

For more information, refer to the `bpdgclone` man page.

For assistance with `vxprint` and other Volume Manager commands, refer to the *Veritas Volume Manager Administrator's Guide*.

- 3 To verify that the clone has been removed, re-enter `vxdg list`.

Sample output:

NAME	STATE	ID
rootdg	enabled	983299491.1025.turnip
VolMgr	enabled	995995264.8366.turnip
wil_test	enabled	983815798.1417.turnip

The clone no longer appears in the list.

Alternate client restore and backup from a snapshot fails

An alternate restore from a snapshot fails when the Veritas Volume Manager (VxVM) version is not the same on both the client and the alternate client. The different version of VxVM also causes a backup from a snapshot operation failure with error 4213.

To perform a successful alternate client restore and a Backup from a Snapshot operation, confirm that the VxVM version is the same on the client and the alternate client.

Upgrade the NetBackup clients to the same version of VxVM.

Restore from a snapshot fails with status 2800

Restore from a snapshot fails with status code 2800 and displays the 'Standard policy restore error' message. The progress log shows "no files matched in the given data range". The restore fails if you select a different path other than what has been mentioned in the backup selection.

For example, suppose `vol6` and `vol7` are volumes mounted on `/mnt/vol6` and `/mnt/vol7` respectively. These mount points are specified in the backup selection. During a restore if you select only `/mnt`, which is the parent directory of the path mentioned for backup selection, the restore fails with status code 2800.

For a successful restore from the snapshot copy, you must select the original path mentioned in the Backup Selections tab, that is `/mnt/vol6` and `/mnt/vol7` or the sub-directory or file.

Raw Partition restore fails with the message 'FlashBackup-Windows policy restore error'

When a raw partition is set for a restore from a snapshot copy, the restore job fails with the error - FlashBackup-Windows policy restore error.

Only a single file restore from a snapshot copy is supported. A raw partition restore from a snapshot copy for a FlashBackup-Windows policy is not supported.

When you need to restore a raw partition, use the storage unit copy (2nd copy) of a snapshot. You must set the storage unit copy as the primary copy.

Snapshot creation fails with error 156

Snapshot creation fails when multiple volumes are mounted to the same mount point. Consider a scenario where `/vol/gvol1` and `/vol/gvol2` are mounted on `/mnt`. Here the volumes `gvol1` and `gvol2` are both mounted on the same mount point `/mnt`. When a backup policy is run to take snapshot backups, the snapshot creation fails with error 156 as both the volumes are mounted on the same mount point.

The reason for snapshot creation failure in the above case is that there would be multiple entries in OS mount table for the same mount point, which is an unsupported configuration in NetBackup.

Veritas recommends that you mount each volume on separate mount points.

Snapshot creation can fail due to multiple reasons. The reason mentioned above is just one of them.

Snapshot fails with error 20

Snapshot creation fails when a mount point used in the backup selection contains a symbolic link. The reason for snapshot failure is that NetBackup does not support a symbolic link as a mount point in the backup selection for a policy.

For example, consider two directory structures `/dir1/dir2` and `/dir3/dir4` and `ln -s /dir1/dir2 /dir3/dir4` is the symbolic link.

If the path `/dir1/dir2/dir4` is used as mount point in backup selection, policy validation skips such mount point and succeeds. Further if user runs manual backup for a policy, Snapshot fails with error 20.

Restore from a snapshot fails with status 133

Restore from snapshot fails with status code 133 and displays the *Invalid request* message.

The restore fails, if you select a path other than the path mentioned in the backup selection.

For example, say the backup selection contains `/ifs/voll/parent/dir1`. During a restore if you select only `/ifs/voll/parent`, which is the parent directory of the path mentioned for backup selection, the restore fails with status code 133.

For a successful restore from the snapshot copy, you must select the original path mentioned in the **Backup Selections** tab, that is `/ifs/voll/parent/dir1` or the sub-directory or file inside the backup selection.

Backup from snapshot fails with error 50

This error occurs when the NetBackup Client and NetBackup Legacy Network services are not restarted properly after configuration for the domain user.

If you are using master or media as backup host then follow these steps to troubleshoot:

- 1 Stop all NetBackup services using the `bpdown.exe`.
- 2 Logon to the NetBackup Client and NetBackup Legacy Network services as the domain user.

But, do not start these services immediate after logon.
- 3 Start all the services together using the `bpup.exe`.

Backup from snapshot parent job fails with error 4213: Snapshot import failed

Job details shows an error like:

"Snapshot export failed. Failed to export share: data_lif is not online. Please check data_lif status on vserver: VSERVER_1."

where, VSERVER_1 is the vserver that is offline.

For a NAS-Data-Protection policy, all the vservers are listed in the client's section of the policy, irrespective of their state. So, you are able to include backup selection from offline SVM, and policy validation succeeds. However, at the time of

Snapshot job fails and the snapshot command does not recognize the volume name

backup-from-snapshot, export operation for those shares fails if the corresponding vserver is offline.

To overcome this error, check the status of the vserver and whether that vserver is reachable. Whenever client and vserver connection is established, SLP retry is successful.

Snapshot job fails and the snapshot command does not recognize the volume name

A snapshot job fails if the volume name exceeds 15 characters.

When you create and name a volume, a prefix or a suffix is added to the volume name. If the volume name contains more than 15 characters, addition of prefix or suffix may make the volume name exceed the limit of 27 characters. When you run the `vxassist snapshot`, command, it does not recognize the lengthy snapshot volume name and the snapshot job fails.

For example, if the primary volume name is **PFltest123456789vol** and the suffix **00043c8aaa** is added to it, the volume name exceeds the limit. The command `vxassist snapshot` does not recognize the name **PFltest123456789vol_00043c8aaa** and the snapshot job fails.

Veritas recommended that you limit the primary volume names to up to 15 characters to create the VxVM mirror snapshots.

Snapshot creation fails with error 4220

Snapshot creation fails due to a zoning issue. If the zoning is not properly done, the target devices/LUNs do not respond well.

For a successful snashot creation, verify that zoning is properly done for the host and the FC/scsi adapter is responding well.

For example, for AIX, you can verify zoning, using the following steps:

To verify if the zoning is correct

- 1 Check the FC/scscli adapter on the host in `/usr/sbin/lsdev -C -c adapter -Sa -F 'name'`

The Output should be:

fcs0

fcs1

- 2 Check FC SCSI I/O Controller in `/usr/lib/methods/cfggefc -l fcs0.`

The Output should be

Output:

- 3 Check time taken to rescan the devices with the FC SCSI I/O Controller in time `/usr/lib/methods/cfggefscsi -l fscsi0.`

The Output should be:

<disk list>

real 2m2.123s the real wall clock time to execute the command

user 0m0.047s the CPU cycle time for the command in user mode.

sys 0m0.024s the CPU cycle time for the command in kernel mode.

Note: If the total time, is greater than 155 sec the snapshot will fail with error 4220.

Snapshot creation fails when the same volume is mounted on multiple mount points of the same host

Snapshot creation fails when the same volume is mounted on multiple mount points of the same host.

For example, when the volume `f3170-7-15:/vol/sample1` is mounted on the mount points `/sample1` on `f3170-7-15:/vol/sample1`

`rsize=32768,wsiz=32768,NFSv3,dev=4000033` and `/test1` on `f3170-7-15:/vol/sample1` `rsize=32768,wsiz=32768,NFSv3,dev=4000034` snapshot creation fails with the following error.

```
mount: f3170-7-15:/vol/sample1 is not mounted on /test1
```

The backup of NFS share mounted by two different mount points for OST_FIM is not supported in this release.

Snapshot-based backup and restore failure

Snapshot-based backups and restores fail if the backup selection that is listed in the NetBackup policy contains any spaces either in the mount points or mount devices. For example,

No spaces in mount points

Block device example: /dev/dg/vol is mounted on /mnt point

NFS example: Filer:/vol/datavol is mounted on /nfs mnt

No spaces in mount devices

Block device example: /dev/dg/vol data is mounted on /mntpoint

NFS example: Filer:/vol/data vol is mounted on /nfsmnt

Multiple snapshot jobs fail with code 156 or 1541.

Multiple snapshot jobs that were started at a high frequency fail with code 156 or 1541.

These errors may occur in a situation when an administrator manually (or by using a script), starts multiple snapshot jobs at a high frequency. (For example, one snapshot job every 5 seconds.)

At the same time, multiple rotation processes begin. The processes operate on the same catalog information, which includes information about existing snapshots. Because the processes work on the same information at the same time, a problem of inconsistency can occur. Some of the processes delete the snapshots and update the catalog while other processes continue to refer to the obsolete information. The result is that the snapshot jobs can end with status codes 156 (snapshot error encountered), 1541 (snapshot creation failed), or other unpredictable errors.

This behavior does not occur for scheduled snapshot jobs, as NetBackup controls the job execution.

FlashBackup policy fails, with multiple backup selections [Cache =]

A FlashBackup policy created with the VxFS_Snapshot fails when multiple backup selections are specified in the same policy.

The VxFS_Snapshot can be used to backup a single file system only. If multiple file systems are backed up using the same policy, the backup fails.

Make sure that you create a separate policy for each file system.

Partial backup failure with 'Snapshot encountered error 156'

There can be a partial backup failure when the NetBackup client and the alternate client that are specified in the policy are located on the same host. The off-host backup functionality ideally shifts the load of backup processing onto an alternate client. The alternate client sends the client's data to the storage device.

When both the client and alternate client are on the same host, the off-host backup purpose is lost. The client load is increases and backups fail.

Note: The client and the alternate client on the same host is not supported.

Backup of file system validation fails with error 223

A Hitachi hardware snapshot backup fails with status 223 and error message 'Snapshot could not be created with ShadowImage method'.

Check the Storage Foundation version, it must be 5.0MP3 RP3 HF9 or later. The Red Hat Linux 4 platform supports only 5.0MP3 RP3 HF9 or later versions.

Policy validation fails if the specified CIFS share path contains a forward slash

Policy validation fails when CIFS share the path specified in backup selection fails as &newProductName; backup does not support the path that contains a forward slash. For example, the path \\NASFiler1\dataShare1 for OST_FIM and C:\backup\testdir for VSS FIM are not supported.

Veritas recommends that you use the back slash when you specify the backup selection. For example, \\NASFiler1\dataShare1 and C:\backup\testdir are valid paths.

Policy validation fails with status code 223

Policy validation fails with status code 223, and displays the error message *An invalid entry was encountered*.

For NASDP policy, NetBackup does not support spaces in name of the NAS Array Head. If the name of NAS Array Head contains spaces, then policy validation would fail with this error.

Avoid spaces in the NAS Array Head name.

An NDMP snapshot policy for wildcard backup fails with error 4201

An NDMP policy for wildcard backup fails when the policy is run after the image expires. In this scenario, the policy has a specific retention period after which the image expires. If a backup is taken immediately after the specified retention period the backup automatically fails. The backup fails due to the conflict in snapshot creation as the DFM server takes some time to delete the existing snapshot-related information.

For a successful backup, run the backup job after some time.

Troubleshooting with bpfis log

The NetBackup `bpfis` log is a good source of troubleshooting information.

See [“UNIX logging directories for backup”](#) on page 178.

See [“Windows logging folders for backup”](#) on page 180.

Limitations of using HP-UX 11.31

HP-UX 11.31 does allow a new device to be present on the same SCSI path on which different device was visible to the host. During the snapshot process, when the old snapshot is deleted and a new snapshot is created, the new snapshot appears on the same SCSI path as the older snapshot. That causes a conflict within the HP-UX system and it logs an error message.

During a snapshot with NetBackup installed on a computer that has HP-UX 11iv3 installed, the Syslog error messages are similar to the following:

```
class : lunpath, instance 15
Evpd inquiry page 83h/80h failed or the current page 83h/80h
```

While configuring a policy for snapshot operations, the NetBackup Administration Console hangs

```
data do not match the previous known page 83h/80h data on
LUN id 0x0 probed beneath the target path (class = tgtpath,
instance = 4) The lun path is (class = lunpath, instance 15).
Run 'scsimgr replace_wwid' command to validate the change
class : lunpath, instance 15
Evpd inquiry page 83h/80h failed or the current page 83h/80h
data do not match the previous known page 83h/80h data on
LUN id 0x0 probed beneath the target path (class = tgtpath,
instance = 4) The lun path is (class = lunpath, instance 15).
Run 'scsimgr replace_wwid' command to validate the change
class : lunpath, instance 15
An attempt to probe existing LUN id 0x4007000000000000 failed
with errno of 14.
0/3/1/0.0x50001fe150070028.0x4007000000000000 eslpt
0/3/1/0.1.27.0.0.0.7 sdisk
64000/0xfa00/0x69 esdisk
```

The administrators of the HP-UX 11iv3 host machines should ignore the log messages if they encounter them during backups with NetBackup.

While configuring a policy for snapshot operations, the NetBackup Administration Console hangs

For on-premises devices, while configuring a policy for snapshot operations, the NetBackup Administration Console hangs.

This issue might arise if the source devices are not responsive or are in a hung state.

Accelerator enabled incremental backup of NetApp NAS volume

Accelerator enabled NAS-Data-Protection policy backups complete volume instead of only the incremental data. This also affects the run optimization.

This issue occurs under the following conditions:

- The policy type is NAS-Data-Protection.
- In the policy's Snapshot options, the value of Access Protocol is Default or NFS3.

- Backup selection has NetApp NAS volumes.

The Accelerator technology optimizes a backup by sending only changed blocks over a network for backup. A two-step process is used to identify the changed files and changed blocks in these files. File attributes and index node (inode) are the key parameters to identify a change. If the files are accessed over NFS version 3, a file on NetApp NAS volume behaves different because of the inode numbers. Same file has different inode numbers across snapshots of the volume if accessed over NFS3. All schedules of backup are based on the snapshot that is created for the run of the policy. A new snapshot with different inode numbers than the previous ones makes accelerator to identify these files as new files. Because of this issue, all files are backed up instead of incremental data only.

To resolve this issue, avoid using NFS version 3 to access the snapshot for accelerator-enabled backups. You can change the Access Protocol to NFS4 for the affected policy. For more details, refer to the [NetApp documentation](#).

Snapshot method: Auto

Error scenario 1: Policy validation fails, after a primary server upgrade, if you create a policy with VSO FIM for older clients and select Snapshot Method as Auto in the NetBackup 10.0 UI.

Error scenario 2: Snapshot jobs fail, if you configure DNAS policy with backup host pool containing older version backup hosts and select Snapshot Method as Auto in the NetBackup 10.0 UI.

The Snapshot Method, Auto is supported only in NetBackup 10.0 onwards. If your environment contains older version backup hosts, select another snapshot method.

Managing nbu_snap (Solaris)

This appendix includes the following topics:

- [About managing nbu_snap](#)

About managing nbu_snap

This topic describes commands for managing snapshots that were made with the nbu_snap method. This topic applies to Solaris systems only.

Cache for nbu_snap

NetBackup uses the nbu_snap method to create a copy-on-write snapshot in the following circumstances (Solaris systems only):

- If the **nbu_snap** method is configured for the back up on the **Snapshot Options** dialog.
- If NetBackup's auto-selection mechanism chose the **nbu_snap** method when the backup was initiated.
- If the client is in a FlashBackup policy and no snapshot method is configured (**Perform snapshot backups** on the Policy display is not selected).

In all cases, a raw partition must be specified as cache for the snapshot. Any number of concurrent nbu_snap backups can use the same cache. However, the cache must be large enough to hold copies of all file system blocks that were changed by user activity while the snapshots were active.

In the NetBackup Administration Console, the raw partition cache can be specified in any of three places.

See [“Entering the cache”](#) on page 97.

About determining cache size

The required size of the cache partition depends on how much user write activity occurs while the snapshot is active. The required cache size does not depend on the size of the file system. The more user activity, the larger the cache must be. You can use the `snaplist` and `snapcachelist` commands to determine the required size of the cache.

See [“Determining a size for the cache partition”](#) on page 95.

If a cache is too small and overflows, all snapshots using the cache become unreadable and the backups that read the snapshots fail.

About terminating nbu_snap

NetBackup ordinarily removes snapshots after the `&newProductName;` backup completes. However, as a result of some system failures, such as a system crash or abnormal backup termination, the snapshot may not have been removed.

Use the `snapoff` command to terminate an `nbu_snap` snapshot that the backup job did not terminate.

See [“snapoff command”](#) on page 210.

See [“Identifying and removing a left-over snapshot”](#) on page 188.

nbu_snap commands

The following commands relate to the `nbu_snap` snapshot method.

snapon command

`snapon` starts an `nbu_snap` snapshot (copy-on-write).

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snapon snapshot_source cache
```

where:

- *snapshot_source* is the partition on which the client's file system (the file system to be "snapped") is mounted
- *cache* is the raw partition to be used as copy-on-write cache.

Example 1:

```
/usr/opensv/netbackup/bin/driver/snapon /var /dev/rdisk/c2t0d0s3
```

Example 2:

```
/usr/opensv/netbackup/bin/driver/snapon /dev/vx/rdisk/omo/tcp1  
/dev/vx/rdisk/omo/sncache
```

The snapshot is created on disk, and remains active until it is removed with the `snapoff` command or the system is restarted.

snaplist command

`snaplist` shows the amount of client write activity that occurred during an `nbu_snap` snapshot. Information is displayed for all snapshots that are currently active.

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snaplist
```

Information is displayed in the following form:

id	ident	size	cached	minblk	err	time
9	6730	2097152	2560	0	0	05/29/03 07:52:18
	device = /dev/vx/rdisk/omaha/tcp1					
	cache = /dev/vx/rdisk/omaha/sncache					
16	7857	2097152	128	0	0	05/29/03 12:16:00
	device = /dev/vx/rdisk/omaha/vol01					
	cache = /dev/vx/rdisk/omaha/sncache					
17	7908	20971520	4224	0	0	05/29/03 12:17:38
	device = /dev/vx/rdisk/omaha/vol03					
	cache = /dev/vx/rdisk/zetab/cache					

If no snapshots are currently active, no output is returned.

Description of output:

<code>id</code>	The snapshot ID. The <code>snapoff</code> command uses this ID to terminate the snapshot.
<code>ident</code>	A unique numeric identifier of the snapshot. <code>ident</code> is the pid of the process that created the snapshot.

size	The size of the client's snapshot source in 512-byte blocks. The snapshot source is the partition on which the client's file system (the file system being backed up) is mounted. Note: <code>size</code> is not a reliable guide to the size of the cache that is needed for the snapshot. The user write activity during the snapshot is what determines the size of the cache needed. See the <code>cached</code> column of this output.
cached	The number of 512-byte blocks in the client file system that were changed by user activity while the snapshot was active. Before being changed, these blocks were copied to the cache partition. The more blocks that are cached as a result of user activity, the larger the cache partition required. However, additional overhead—which is not shown in this <code>cached</code> value—is required in the cache. To see the total space that is used in a particular cache partition, use the <code>snapcachelist</code> command.
minblk	In the partition on which the file system is mounted, <code>minblk</code> shows: the lowest numbered block that is monitored for write activity while the snapshot is active. Only FlashBackup policies use <code>minblk</code>.
err	An error code; 0 indicates no error. If a snapshot has encountered an error, the <code>err</code> is non-zero and the snapshot is inaccessible. It can be terminated using <code>snappoff</code> and the snapshot ID. Error codes are identified in <code>/usr/include/sys/errno.h</code>. Also, error messages may be found in <code>/var/adm/messages</code>.
time	The time at which the snapshot was started.
device	The raw partition containing the client's file system data to back up (snapshot source).
cache	The raw partition used as cache by the copy-on-write snapshot process. Make sure that this partition is large enough to store all the blocks likely to be changed by user activity during the backup. To determine the total space that is used in a particular cache partition by all active snapshots, use the <code>snapcachelist</code> command.

snapcachelist command

`snapcachelist` displays information about all partitions currently in use as `nbu_snap` caches. This command shows the extent to which the caches are full.

Note: snaplist and snapcachelist can also be used to monitor an nbu_snap snapshot that a NetBackup policy started. Note that once the backup completes, NetBackup removes the snapshot. As a result, the snaplist and snapcachelist commands no longer return any output.

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snapcachelist
```

If no snapshots are currently active, no output is returned.

Output is of the form:

device	free	busy
/dev/rdisk/c0t4d0s0	238528	264472

Description of output:

device	The raw partition being used as cache.
free	The number of 512-byte blocks unused in the cache partition.
busy	The number of 512-byte blocks in the client data that changed while the snapshot was active. Before being changed, these blocks were copied to the cache partition by the nbu_snap copy-on-write process. For each cache device that is listed, busy shows the total space that was used in the cache. You can use this value as a sizing guide when setting up raw partitions for nbu_snap cache. When a cache is full, any additional change to the client data causes the copy-on-write to fail: the snapshot is no longer readable or writable. Reads or writes to the client data continue (that is, user activity is unaffected). The failed snapshot, however, is not terminated automatically and must be terminated using snapoff.

snapstat command

snapstat displays diagnostic information about the snap driver.

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snapstat
```

snapoff command

snapoff terminates an nbu_snap snapshot.

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snapoff snap1 ... snapn
```

where *snapx* is the numeric id of each copy-on-write process to be terminated. (Use `snaplist` to display the id of active snapshots.)

If `snapoff` is successful, a message of the following form is displayed:

```
snapshot 1 disabled  
snapshot 2 disabled  
...  
snapshot n disabled
```

If `snapoff` fails, an explanatory message is displayed. Error codes are identified in `/usr/include/sys/errno.h`.

Warning: Do not terminate a snapshot while the backup is active: corruption of the backup image may result.

Overview of snapshot operations

This appendix includes the following topics:

- [Introduction to snapshot operations](#)
- [Pre and post snapshot creation operations](#)
- [About quiescing the system](#)
- [About quiescing the database application](#)
- [About quiescing the stack](#)
- [File system quiesce](#)
- [Volume manager data caching](#)
- [How copy-on-write works](#)

Introduction to snapshot operations

This topic presents background information on snapshot operations and copy-on-write snapshots.

Before a snapshot is created, there are certain operations that NetBackup performs as a part of the pre- and post snapshot creation process. The snapshots that are created after these operations contain authentic data. Copy-on-write (COW) is a snapshot type that gives a detailed account of data as it existed at a certain moment.

For a diagram of the daemons or services involved in creating and backing up a snapshot, refer to the [NetBackup Troubleshooting Guide](#).

Pre and post snapshot creation operations

NetBackup performs several vital functions before snapshot creation. Without the pre-processing, the integrity of the snapshot cannot be guaranteed and the backup data may be of no value.

Table B-1 Pre and post snapshot creation operations

Step	Action
Step 1	Backup process requests database quiesce.
Step 2	Database application quiesces (must wait for transactions to complete).
Step 3	Lock and flush the file system.
Step 4	Create the snapshot.
Step 5	Unlock the file system.
Step 6	Release (unquiesce) the application.
Step 7	Back up the snapshot.
Step 8	(Optional:) Remove the snapshot.

Note: Steps 1, 2, and 6 in [Table B-1](#) apply only to databases, such as those requiring NetBackup for Oracle &newProductName;.

About quiescing the system

Before a useful snapshot can be created, the data to back up must be transactionally consistent or complete. A transaction is a single data action, such as updating a patient's record in a medical database, or creating a record for a new patient. Such a transaction is composed of multiple I/O requests (search, copy, send, write, and so forth). Until the transaction's I/O requests are complete, the data is inconsistent and may be unsuitable for backup.

Transactions affect all levels of the storage management stack (file system, volume manager, and so forth). A transaction generates further transactions as a request is handed off to the next level of the stack. For instance, in the file system, an I/O request from a database application constitutes a transaction and may be split into many disk references. All these disk references must be complete for the original request to be fulfilled. Thus, the creation of the snapshot must be coordinated with any application or process that can affect the transactional consistency of the data.

The means of coordination is called quiesce (literally, to make quiet or place in repose). Quiesce involves pausing the database application or process until the data is transactionally consistent. Applications and the storage management stack must all be quiesced before a useful snapshot can be made.


About quiescing the database application

Most database applications are transactionally consistent only at particular points in time. Sometimes, they are consistent only after they have been shut down. Since many database applications must be available constantly, many applications can reach transactional consistency at regular intervals or in response to an external event. This process is called application quiesce, described in this section.

In database application quiesce, an external signal or message is sent to a receptive database. In response, the database finishes the current transaction or group of transactions and tells the snapshot consumer when the transactions are complete. The database then waits for the indication that normal operations can resume. After the database indicates that it has reached a state of transactional consistency, the final steps of creating the snapshot can proceed.

Once the snapshot has been created, another signal is sent to the waiting database to resume normal operations. This procedure is called unquiescing the application.

Figure B-1 Dialog for quiesce or unquiesce



About quiescing the stack

The storage management stack is a layered arrangement of software elements. An I/O request from a database application passes from element to element until a hardware request to move data reaches the storage network. Each stack element performs a variety of functions, some of which treat I/O requests like transactions to assure their completion. Before a snapshot is created, the stack must be quiesced (made transactionally consistent).

Since the file system is the front-line interface to applications for managing files and I/O, file system quiesce is critical to quiescing the stack.

File system quiesce

Quiescing the file system consists of two major tasks as listed in the following points:

- Prohibit new I/O requests from initiating, which is called locking the file system.
- Flush file system cache (write cached data back to disk). The system must complete any outstanding application I/O and note completion of outstanding metadata updates.

Volume manager data caching

As in a file system, the volume manager's data caching may have to be flushed and disabled until the snapshot is created. If volume manager caching is enabled, data that is required for a consistent image may linger in volume manager cache rather than residing on disk when the snapshot is created.

How copy-on-write works


A copy-on-write snapshot is a detailed account of data as it existed at a certain moment. Unlike a mirror, a copy-on-write is not a copy of the data, but a specialized account of it.

The copy-on-write process works as follows: when a snapshot is required, any unfinished transactions or changes to the source data are allowed to complete, but new changes are temporarily stalled. The source is momentarily idled (made quiescent). Once the copy-on-write is activated, new transactions or changes (writes) to the source data are allowed to take place. However, the copy-on-write process briefly intercepts or holds the first write request that is issued for any particular block of data. While it holds those requests, it copies to cache the blocks affected by those writes, and keeps a record of the cached blocks. In other words, it reads each

source block that is about to change for the first time. Then it copies the block's current data to cache, and records the location and identity of the cached blocks. Then the intercepted writes are allowed to take place in the source blocks.

Figure B-2 shows the copy-on-write process.

Figure B-2 Copy-on-write process



The following table lists the phases that have been depicted in the diagram:

Phase	Action
Phase 1	Image of source data is frozen; copy-on-write is activated.
Phase 2	New write requests to s4, s7, s8 are held by copy-on-write process (see arrows).


Phase	Action
Phase 3	Copy-on-write process writes contents of blocks s4, s7, and s8 to cache. These blocks write to cache only once, no matter how many times they change in the source during the snapshot.
Phase 4	Copy-on-write process keeps a record of the number of writes to cache.
Phase 5	Write requests are now allowed to take place.

The immediate results of the copy-on-write are the following: a cached copy of the source blocks that were about to change (phase 3), and a record of where those cached blocks are stored (phase 4).

The copy-on-write does not produce a copy of the source. It creates cached copies of the blocks that have changed and a record of their location. The backup process refers to the source data or cached data as directed by the copy-on-write process.

[Figure B-3](#) shows the process for backing up a copy-on-write snapshot.

Figure B-3 Backing up a copy-on-write



The following table lists the phases that have been depicted in the diagram:

Phase	Action
Phase 1	Backup reads source data from s0, s1, s2, s3
Phase 2	At s4, copy-on-write tells backup to read c0 instead of s4
Phase 3	Next, the backup reads s5 and s6 from the source.
Phase 4	At s7 and s8, copy-on-write tells backup to read c1, c2 instead of s7, s8.
Phase 5	Backup continues reading of the source or cache, as directed by copy-on-write.
Phase 6	When backup completes, backup data is identical to original source.

As this diagram shows, an accurate backup image is obtained by combining the unchanged portions of the data with the cache. When a backup of the snapshot begins, the backup application copies the source data (phase 1) until it encounters a block that changed after the copy-on-write process started. The copy-on-write tells the backup to skip that changed block and read in its place the cached (original) copy (phase 2). The backup application continues copying source data (phase 3) until it comes to another changed block. Cache is read again (phase 4) as the copy-on-write process dictates. The backup, when finished, is an exact copy of the source as it existed the moment the copy-on-write was activated.

Index

Symbols

156 status code 184
3pc.conf file 32

A

abnormal termination 188, 207
Active Directory 72, 110
 and FlashBackup 72
actual device file name 76–77
ALL_LOCAL_DRIVES entry 28, 42, 153
Allow multiple data streams 50
alternate client
 defined 29
Alternate client backup
 configuring 47
alternate client backup 14, 20, 75
 and FlashSnap 103
 and split mirror 22
 and VVR 106
 introduction 21
 requirements 47
 restrictions 48
 testing setup 103, 107, 109
APP_NBU_VVR 107
arbitrated loop 30
archives 156
auto snapshot selection 48
automatic backup 155

B

Backup
 Archive
 Restore 157
backup
 agent 14, 29
 automatic 155
 local 18
 logs 178, 180
 manual 156

backup (*continued*)
 off-host
 configuration 41–42
 prerequisites 152
 raw partition 153
 retention period
 NAS_Snapshot 108
 scripts 46
 techniques (overview) 15
 types supported 76
 user-directed 156
backup agent 42
backup retention level 87
Backup Selections list 41
 ALL_LOCAL_DRIVES entry 28, 42, 153
 and Instant Recovery 88
 block vs. character device 153
 directives 80
 FlashBackup 77
 symbolic link 49
block device file (vs character) 153
bp.conf file 178
bpbkar
 log 178
 process 187
bpbrm log 178–179
bpdgclone command 195
bpfis
 command 166, 169, 188
 log 178
bppfi log on client 178
bprd log 179
bprestore log 179
bpvm log 178–179
BusinessCopy method 185

C

cache
 definition of 29
 flushing for quiesce 215
 object 101

- cache (*continued*)
 - overflow 207
 - partition 80
 - partition out of space 186
 - requirements 94
 - size 207
 - specifying raw partition for 94
- CACHE= directive 78, 80, 97
- Change Policy dialog 44, 46
- character device file (vs block) 94, 153
- checking logs 177
- checkpoint
 - removing (VxFS) 190
 - VxFS file system 98
- CLARiiON
 - fractured (split) clone after rollback 137
- client data
 - prerequisites for off-host backup 152
- client list 77
- client read timeout 154
- clone
 - fractured (split) after rollback 137
 - removing 190, 193
 - VxVM disk group 193
- Cluster Volume Manager. *See* CVM
- compression 39
- configuration
 - backup selections 77
 - client list 77
 - prerequisites 38
 - procedure 40, 43, 87
 - snapshot parameters 45
 - supported data types 72
- Configure snapshot policy 39
- copy manager (see third-party copy) 31
- copy-on-write
 - defined 29
 - overview 16, 215
- cp command 165
- cross mount points (disabled) 50
- customer support 183
- CVM
 - about support for 111
 - enabling vxvm or FlashSnap snapshots in 112
 - execute VxVM commands remotely 113

D

- data consistency (quiesce) 213

- data mover 32
 - Network Attached Storage 43
- data mover (see third-party copy) 31
- database files (Windows) 72, 110
- DCO plexes 104
- deleting
 - a clone 193
- deport
 - disk group 103, 105
 - snapshot 167, 170
- device
 - serialization 153, 185
- Device Configuration wizard 38
- device file name (FlashBackup) 77
- directives
 - for scripts 46
- directives (Backup Selections list) 80
- directories for troubleshooting 177
- disabling snapshots 50
- disk
 - restoring snapshot from 165
 - SCSI vs. IDE 184
 - visibility 184
- disk array methods
 - Hitachi 127
 - HP-XP 132
 - IBM DS4000 122
 - IBM DS6000 and DS8000 116
 - troubleshooting 135
 - verify NetBackup access to arrays 115
- disk array snapshot methods
 - dynamic multipathing 114
 - HBAs 115
 - OS-specific configuration 114
 - persistent target bindings 115
- disk group
 - clone 193
 - shared in VxVM 94, 103
- disk restore
 - caution 168
- disk snapshot
 - restoring from 164–165
- disk storage units 153
- DSCLI 116
- duplex mode and performance 184

E

- encryption 39
- exceptions to exclude list 72

exclude list 72, 77

Extended Copy command 14, 31, 43

extent

defined 30

F

fast mirror resynch 102

FastResynch 104

Fibre Channel

defined 30

types supported 30

file pathname (max length) 42, 49

file promotion 15

file system

multi-volume 98, 101

file systems

defined 30

quiescing 215

files list (see Backup Selections list) 41

FlashBackup 206

actual device name 77

and VxFS_Checkpoint 98

backup selections 77

device file name 76

features 71

files to back up 76–77

restoring 156

FlashSnap method

and status 156 184

deporting the disk 103

preparing for alternate client backup 103

restoring from disk image 167, 170

with VxVM shared disk group 103

flush file system 215

FORCE_RESTORE_MEDIA_SERVER option 163

fsckptadm command 190

full-sized instant snapshots 102–103

G

GID 47

glossary of terms 29

H

hardware-level disk restore 168

Hitachi disk arrays 127

configure NetBackup for 128

configuring a NetBackup policy for 130

nbfireshcan 128

Hitachi disk arrays (*continued*)

obtain identifiers 129

pair status 128

RAID Manager version 127

software requirements 127

Hitachi_CopyOnWrite

configuring a NetBackup policy for 130

Hitachi_ShadowImage

configuring a NetBackup policy for 130

HP 80

HP-XP disk arrays 132

configure NetBackup 133

nbfireshcan 133

obtain identifiers 133

pair status 132

software requirements 132

HP_XP disk arrays

configuring a NetBackup policy 134

HP_XP_BusinessCopy

configuring a NetBackup policy 134

HP_XP_Snapshot

configuring a NetBackup policy 134

I

I/O components

supported 72

IBC messages 107

IBM DS4000 disk array 122

configure a NetBackup policy for 126

configure for NetBackup 125

nbfireshcan 124

SMclient 123

software requirements 123

verify NetBackup client access to 124

IBM DS6000 and DS8000 disk arrays 116

configure array for NetBackup 117

configure NetBackup policy for 121

IBM_DiskStorage_FlashCopy 121

software requirements 116

IBM_DiskStorage_FlashCopy 121

IBM_StorageManager_FlashCopy 122

configure a NetBackup policy for 126

IDE vs. SCSI 184

importing disk group 105–106

include list 77

inquiry page code 83 153, 185

Instant Recovery

defined 30

Schedule tab 88

Instant Recovery (*continued*)
 selecting in policy 87
 snapshot only 87
 volume name restriction 101
 instant snapshots 101–102

J

jbpSA 157

K

Keep snapshot after backup 188
 restoring from image 164–165
 kernel messages 179

L

left over snapshot 188, 207
 Limit jobs per policy 50
 limitations 28
 links (in Backup Selections list) 49
 local host backup method
 network configuration for 18
 lock file system 215
 logging
 directories to create 177
 VxMS 181
 logical volume (as raw partition) 94
 logs 177–180
 creating for UNIX 178
 creating for Windows 179
 loop (Fibre Channel) 30
 LVM 72

M

manual backup 156
 mapping
 defined 30
 Maximum jobs per client 50
 Maximum multiplexing per drive 50
 maximum pathname length 42, 49
 Maximum Snapshots (Instant Recovery only) 45, 87, 108
 Media multiplexing 50
 media server (see NetBackup Media Server) 42
 messages file 179
 method
 selecting off-host backup 41–42
 selecting snapshot 44

mirror 16
 compared to copy-on-write 18
 defined 30
 fast resynch 102
 overview 16
 VxVM snapshot 101
 mklogdir script 178
 mklogdir.bat 179
 mover.conf file 32
 multi-volume system (VxFS) 98, 101
 multiple data streams 50
 configuring 50
 multiplexing 153
 MVS 98, 101

N

NAS
 off-host backup 20
 NAS filer
 as backup agent 43
 NAS_Snapshot 88, 108
 backup retention period 108
 logging info 178, 180
 nbfirescan 115
 nbu_snap method 93, 206
 with VxVM shared disk group 94
 NDMP host
 as backup agent 43
 NDMP snapshot 108
 ndmp unified log (VxUL) 178, 180
 NetBackup Media Server 20, 31
 and storage units 153
 network diagram of 27
 selecting 42
 Network Attached Storage 20
 Network Attached Storage (data mover) 43
 network interface cards 184
 NEW_STREAM directive 80
 NIC cards and full duplex 184

O

off-host backup 41–42
 and multiplexing 153
 NAS 20
 overview 19
 prerequisites for 152
 raw partition 153
 type of disk (SCSI vs. IDE) 184

- off-host backup (*continued*)
 - with data mover 32
- Open File Backup
 - disabling 50
 - license 33
- operating system
 - patches 28
- Oracle 28
- overview of snapshot operations 212

P

- page code 83 153, 185
- pairresync command 191
- pairsplit (Hitachi) 128
- pairsplit (HP-XP) 132
- parameters for snapshots 45
- partitions
 - Windows 72
- patches 28, 184
- pathname length 42, 49
- Perform snapshot backups 206
- performance
 - increasing tape 185
- peripherals (latest info on web) 32
- physical device (as raw partition) 94
- platform requirements 28
- platforms supported 34
- policy
 - how to select type of 40
- Policy dialog 40, 74
- policy_name (on mover.conf file) 185
- primary vs. alternate client 21
- promotion
 - file 15
- provider 184

Q

- query snapshot 166, 169, 188
- quiesce 213, 216

R

- RAID 5 101
- raw partition 77, 80
 - block vs. character device 153
 - defined 31
 - not supported with VxFS_Checkpoint 98
 - specifying for cache 94
- recovery procedure 188, 207

- Registry 72, 110
 - and FlashBackup 72
- remote snapshot (see alternate client backup) 21
- removing
 - clone 193
 - snapshots 188, 207
- replicated host 25
- replication
 - for alternate client backup 106
 - testing setup for alternate client backup 107
- Replication Director 20, 186
- requirements for NetBackup 28
- restore 156
 - configurations 163
 - file promotion 15
 - from disk image 164–165
 - hardware-level 168
 - logs 179–180
- Restore everything to its original location 159
- restrictions 28, 72
- resyncfromreplica option 167, 170
- resynchronization of mirror 102
- resynchronize
 - disks 191
- retention level for backup 87
- RMAN 31
- rollback 156
 - causes fractured (split) clone 137
- root
 - specifying as snapshot source 50
- RSM Database 72, 110

S

- SAN 152
 - defined 31
- Schedule tab
 - Instant Recovery 88
- scripts
 - running before/after backup 46
- SCSI Inquiry Page Code 83 153, 185
- SCSI serialization 153, 185
- SCSI vs. IDE 184
- serial numbers 153, 185
- serialization 153, 185
- Shadow Copy Service (see VSS) 48, 108
- ShadowImage method 185
- shared disk group (VxVM) 94, 103
- SIZE_DATA_BUFFERS 185

- snap
 - removing 189
- snappcachelist command 209
- snaptctl 72
 - driver log 179
 - overview 206
- snaptlist command 189, 208
- snaptoff command 189, 210
- snapon command 207
- snapshot
 - auto selection 48
 - back up to local storage 18
 - configuration 43, 87
 - copy-on-write vs mirror (how to choose) 18
 - disabling 50
 - ID 166, 169, 188
 - mirror
 - defined 16
 - mirror (creating) 104–105
 - mirror (VxVM) 101
 - on Solaris client
 - troubleshooting 187
 - overview 15
 - pre-processing for 213
 - removing 188, 207
 - restoring from disk 165–166, 169
 - selecting method of 44
 - source
 - for symbolic link 49
 - state file 36
 - volume (creating) 104–105
- Snapshot Client
 - access web info 32
- Snapshot Options dialog 44, 48
- Snapshots only
 - on Schedule tab 87–88
- snaptstat command 210
- Solaris
 - version requirements 28
- space-optimized snapshots 101–102
 - snapshot methods supported 103
- split mirror backup (alternate client) 22
- stack quiescing 215
- Standard policy
 - restoring 156
- state file 36
- status 13
 - and NetBackup Media Server off-host 154
- status code 156 184

- Storage Checkpoint 72
- storage devices 38
- storage unit
 - restrictions for data movers 153
- storage_unit_name version of mover.conf file 185
- streams
 - allow multiple data 80
- support web site 32
- supported
 - data types 72
 - platforms 34
- switched fabric 30
- symbolic links 49
- symmmir command 191
- synchronize disks 191
- system requirements for NetBackup 28
- System State 72, 110
- system-protected files 110
- system-protected files and FlashBackup 72

T

- tape
 - increasing performance 185
- Terminal Services Database 72, 110
- termination 188, 207
- terminology 29
- third-mirror (vxvm) 101
- third-party copy
 - and multiplexing 153
 - defined 29
 - device configuration 42
- third-party copy device 32
 - and storage units 153
 - defined 31
- Third-Party Copy Device off-host backup 75
- TimeFinder method 185
- troubleshooting 177
 - directories to create 177
 - disk array methods 135
- types of backups supported 76

U

- UFS file system 93
- UID 47
- umount command 167, 191
- unified logging 178, 180
- unmount
 - checkpoint 191

- unmount (*continued*)
 - snapshot 167
- unquiesce 214
- UNSET directive 81–82
- UNSET_ALL directive 81–82
- user-directed
 - archives 156
 - backup 156

V

- VERBOSE setting for logs 178
- Veritas Federated Mapping Services 31
- Veritas Volume Manager 101
- Veritas Volume Manager cluster. *See* CVM
- VMware
 - ALL_LOCAL_DRIVES 42
- volume
 - defined 32
- VSS method 108
- VVR method 26, 107
 - preparing for alternate client backup 106
- vxassist 101, 104–105, 189–190
- vx dg command 104, 106, 167, 170
- vx dg list command 195
- vx disk command 167
- VxFS clone
 - removing 190
- VxFS file system 72, 93
- VxFS multi-volume file system 98, 101
- VxFS_Checkpoint method 72, 98
- VxFS_Snapshot method 100
- vxibc command 107
- VxMS 30–31
- VxMS logging 181
- vxprint 189
- vxprint command 190
- vxrecover command 167
- VxVM
 - and RAID 5 101
 - clone of disk group 193
 - instant snapshots 102
 - mirror 101
 - required version 100
 - shared disk group 94, 103
 - Volume Manager 101
 - volume name restriction 101
- vxvm method 100–101
- vxvol 101, 104

W

- web access to recent Snapshot Client info 32
- whatrmver 127
- wildcards in Backup Selections list 42
- Windows
 - open file backup 33
 - OS partitions 72
 - System database files 72, 110
- Windows Shadow Copy Service 109