

Veritas NetBackup Snapshot Client Administrator's Guide

UNIX, Windows, and Linux

Release 8.0

VERITAS™

Veritas NetBackup™ Snapshot Client Administrator's Guide

Document version: 8.0

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo 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
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

.

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

The latest documentation is available on the Veritas website:

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

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introduction	15
	Snapshot Client features at a glance	15
	Snapshot Client features	17
	About snapshots	17
	About snapshot methods	17
	About snapshot providers	18
	About off-host backup support	19
	About Instant Recovery	20
	About FlashBackup policies	20
	About snapshot methods for disk arrays	20
	About block level incremental backup	21
	About Snapshot Client and NDMP	21
	About snapshot basics	21
	About the copy-on-write snapshot type	22
	About the mirror snapshot type	22
	Benefits of copy-on-write versus mirror	23
	About local backup of a snapshot	24
	Off-host backup overview	25
	About file and volume mapping methods	26
	Off-host backup methods	27
	About alternate client backup	27
	FlashBackup and alternate client combination example	33
	NetBackup media server data mover example (UNIX only)	34
	About the NDMP data mover	35
	Snapshot Client requirements	35
	Snapshot Client restrictions	36
	Snapshot Client terminology	38
	Snapshot Client assistance	42
	About open file backups for Windows	43
Chapter 2	Installation	45
	Installation Prerequisites for Snapshot Client	45
	Snapshot Client installation notes	46
	Adding a Snapshot Client license key on UNIX	47
	Adding a Snapshot Client license key on Windows	47

About distributing Client Software in mixed-platform environments	48
About creating log directories	48
About the snapshot state file	48

Chapter 3 Policy configuration 51

Notes on Snapshot Client policies	51
Configuring a Snapshot Client policy	52
Backup Selections tab options when configuring a policy	55
Off-host backup configuration options	57
Automatic snapshot selection	59
Selecting the snapshot method	60
Snapshot methods	62
Configuration parameters for Snapshot Client	65
Snapshot Resources	71
Configuring backup scripts	71
About using alternate client backup	73
Alternate client backup requirements	73
Configuring alternate client backup	74
Before running the alternate client backup	75
Example alternate client backup configurations	75
Policy configuration tips	76
Maximum pathname length	76
Snapshot tips	76
Multiple data streams	77
About incremental backup of mirror-based snapshots	78
About disabling snapshots	79
Disabling Open File Backups on Windows	79
Disabling Snapshot Client snapshots	79

Chapter 4 FlashBackup configuration 81

About FlashBackup	81
FlashBackup restrictions	82
Restores of Windows encrypted files and hard links	83
Configuring a FlashBackup policy	83
Configuring FlashBackup policy for backward compatibility (UNIX only)	87
About the cache partition	89
Requirements for the cache partition	90
Directives for multiple data streams	90

Chapter 5	Instant Recovery configuration	93
	About Instant Recovery capabilities	93
	Instant Recovery requirements	94
	Instant Recovery restrictions	94
	Giving full server privileges to the media server	96
	About Instant Recovery	97
	About snapshot and backup for Instant Recovery	98
	About NetBackup catalog maintenance	98
	About snapshot management	99
	Means of controlling snapshots	100
	Configuring a policy for Instant Recovery	102
	About sizing the cache for Instant Recovery copy-on-write	
	snapshots	104
	Cache size during restore	105
	Setting an adequate size for the snapshot cache	105
	Large restores from an Instant Recovery snapshot	106
	About configuring VxVM	106
	Creating a snapshot mirror	107
	About creating Instant Snapshots	108
	Using the VxVM 3.5 GUI to configure VxVM mirrors	110
	Modifying the VxVM or FlashSnap resync options for point in time	
	rollback	111
	Instant Recovery for databases	111
	About storage lifecycle policies for snapshots	112
	Configuring a storage lifecycle policy to manage snapshot-based	
	backups for Instant Recovery	112
	Storage lifecycle policies and Snapshot Client	
	troubleshooting	114
Chapter 6	Network Attached Storage (NAS) snapshot configuration	117
	About NAS snapshot overview	117
	Notes on NAS_Snapshot	118
	Logging on to the NetBackup Client Service as the Administrator	119
	Setting up a policy for NAS snapshots	120
	NAS snapshot naming scheme	122
Chapter 7	Configuration of software-based snapshot methods	123
	Software-based snapshot methods	123
	About nbu_snap	123

Cache device requirements	124
About VxFS_Checkpoint	128
About VxFS_Snapshot	130
About VxVM	130
About FlashSnap	133
About VVR	136
About NAS_Snapshot	138
About VSP	138
About VSS	139

Chapter 8	Support for Cluster Volume Manager Environments (CVM)	141
	About support for CVM environments	141
	Note on NetBackup and CVM	142
	About enabling VxVM or FlashSnap snapshots in a CVM environment	142
	About enabling the NetBackup client to execute VxVM commands on the CVM master node	143

Chapter 9	Configuration of snapshot methods for disk arrays	145
	About the new disk array snapshot methods	145
	About array-specific methods vs array-independent methods	146
	Advantages of the new array-specific methods	147
	About types of disk array methods	147
	Important disk array method notes and restrictions	148
	Disk array methods at a glance	149
	Disk array configuration tasks	151
	Configuration tasks for the array administrator	151
	Configuration tasks for the NetBackup administrator	152
	Disk array configuration tasks diagram	153
	OS-specific configuration tasks	154
	About dynamic multi-pathing	154
	HBA configuration	154
	About Solaris sd.conf file	155
	Linux modprobe.conf file	156
	Verifying NetBackup client access, zoning, and LUN masking	156
	About VSS configuration (Windows)	158
	Note on NetBackup array credentials	158
	Initial configuration of certain arrays	158

About EMC CLARiiON arrays	162
EMC CLARiiON software requirements for UNIX	163
Veritas support for VSS Snapshot and EMC CLARiiON	163
Diagram of installed software for EMC CLARiiON	163
Verifying connectivity from client to array	165
About resolving host names on the network	166
Configuring NetBackup to access the CLARiiON array	166
Adding clients to a CLARiiON storage group	167
Configuring for EMC_CLARiiON_SnapView_Clone	168
Creating a clone private LUN with the EMC Navisphere Web interface	169
Creating a clone group and select a LUN as source	169
Adding clone LUNs to the clone group	170
Obtaining the device identifier for each source and clone LUN	172
About configuration for EMC_CLARiiON_SnapView_Snapshot	173
Configuring a reserved LUN pool for the storage processors	174
Configuring a NetBackup policy for a CLARiiON array method	175
Common CLARiiON array configuration problems	176
About EMC Symmetrix arrays	176
EMC Symmetrix DMX software requirements	177
Clone emulation flag can cause snapshots to fail	177
EMC snapshot operation fails	177
Support for EMC Symmetrix with Volume Shadow Copy Service	178
Prerequisites for using EMC Symmetrix disk arrays	178
Configuring NetBackup clients to use EMC Symmetrix	178
About configuring NetBackup to access the Symmetrix array	180
About configuration for EMC_TimeFinder_Mirror	180
About configuration for EMC_TimeFinder_Clone	181
About configuration for EMC_TimeFinder_Snap	182
Configuring a policy for EMC_TimeFinder methods	182
About HP EVA arrays	184
Prerequisites for working with HP EVA arrays	184
HP EVA software requirements for UNIX	184
Diagram of installed software for HP EVA	185
Veritas support for VSS Snapshot and HP EVA	185
Verifying connectivity from clients to array using SSSU 5.0	186
Configuring NetBackup to access the EVA array	188
Configuring a NetBackup policy for an HP EVA array method	189

HP EVA restrictions	189
About IBM DS6000 and DS8000 arrays	190
IBM DS6000 and DS8000 software requirements	190
Preconfiguration for IBM arrays	190
Configuring NetBackup to access the IBM DS6000 or DS8000 array	190
Configuring the IBM array for NetBackup	191
Using DSCLI commands to obtain unique IBM identifiers	192
Configuring a NetBackup policy for IBM_DiskStorage_FlashCopy	195
For further reference on IBM arrays	196
About IBM DS4000 array	196
Array preconfiguration tasks	196
IBM 4000 software requirements	197
Verifying NetBackup client access, zoning, and LUN masking	197
Configuring NetBackup to access the IBM DS4000 array	199
Configuring the IBM 4000 array for NetBackup	199
Configuring a NetBackup policy for IBM_StorageManager_FlashCopy	200
About Hitachi SMS/WMS/AMS, USP/NSC, USP-V/VM	201
Hitachi array software requirements	201
Preconfiguration for Hitachi	202
About communication between NetBackup and the Hitachi array	202
Determining if the Hitachi command devices are visible	202
About configuring the Hitachi array for NetBackup	203
Obtaining the Hitachi array serial number and the unique device identifiers	203
Configuring a NetBackup policy for Hitachi_ShadowImage or Hitachi_CopyOnWrite	204
About HP-XP arrays	206
HP-XP array software requirements	206
Preconfiguration for HP-XP	206
About communication between NetBackup and the HP-XP array	207
Determining if the HP-XP command devices are visible	207
About configuring the HP-XP array for NetBackup	207
Obtaining the array serial number and unique HP-XP identifiers	207
Configuring a NetBackup policy for HP_XP_BusinessCopy and HP_XP_Snapshot	208
About array troubleshooting	209

	Troubleshooting issues pertaining to all arrays	209
	Troubleshooting NetBackup and EMC CLARiiON arrays	210
	Troubleshooting NetBackup and EMC Symmetrix arrays	212
	Troubleshooting NetBackup and HP EVA arrays	213
	Troubleshooting IBM DS6000 and DS8000 arrays	214
	Troubleshooting IBM4000 arrays	217
	Troubleshooting Hitachi arrays	219
Chapter 10	Notes on Media Server and Third-Party Copy methods	227
	Disk requirements for Media Server and Third-Party Copy methods	227
	Directives for Media Server and Third-Party Copy methods	228
	Storage units for Media Server and Third-Party Copy methods	228
	Preventing multiplexing on a third-party copy backup	228
	Raw partition backups	228
	Increasing the client read timeout for all clients	229
	Further information on off-host data mover backups	229
Chapter 11	Backup and restore procedures	231
	About performing a backup	231
	About performing a restore	232
	About restores from a FlashBackup backup	232
	Notes for FlashBackup and UNIX client restore	233
	Notes for FlashBackup and Windows client restore	233
	Restoring a large number of files in a clustered file system (VxFS on UNIX Only)	234
	Instant Recovery restore features	234
	About Instant Recovery: block-level restore	234
	About Instant Recovery: file promotion	235
	About Instant Recovery: Fast File Resync (Windows clients only)	236
	About Instant Recovery: point in time rollback	238
	Notes for restoring individual files from an Instant Recovery snapshot	242
	About configurations for restore	243
	About restoring over the LAN	243
	About restoring over the SAN to a host acting as both client server and media server	244
	About restoring directly from a snapshot	245
	About restoring from a disk snapshot	246
	About restoring on UNIX	246

About restoring on Windows	249
Chapter 12 Troubleshooting	253
About gathering information and checking logs	254
Logging directories for UNIX platforms	254
UNIX logging directories for backup	255
UNIX logging directories for restore	255
snapctl driver messages	256
Logging folders for Windows platforms	256
Windows logging folders for backup	256
Windows logging folders for restore	257
Configuring VxMS logging	257
Customer support contact information	260
Latest patches and updates	260
Snapshot provider information	261
Important notes on Snapshot Client	261
Snapshot Client installation problems	263
FlashBackup and status code 13	263
The FlashBackup cache partition may have run out of space	263
Removing stale snapshots (Solaris)	263
Single file restore from a FlashBackup Instant Recovery snapshot of a file protected by Windows VSS writer fails	264
Identifying and removing a left-over snapshot	264
Removing a VxVM volume clone	270
Alternate client restore and backup from a snapshot fails	272
Restore from a snapshot fails with status 2800	272
Raw Partition restore fails with the message 'FlashBackup-Windows policy restore error'	273
Snapshot creation fails with error 156	273
Snapshot fails with error 20	273
Snapshot job fails and the snapshot command does not recognize the volume name	274
Snapshot creation fails with error 4220	274
Snapshot creation fails when the same volume is mounted on multiple mount points of the same host	275
Snapshot-based backup and restore failure	276
Multiple snapshot jobs fail with code 156 or 1541.	276
FlashBackup policy fails, with multiple backup selections [Cache =]	276
Partial backup failure with 'Snapshot encountered error 156'	277
Backup of file system validation fails with error 223	277

	Policy validation fails if the specified CIFS share path contains a forward slash	277
	An NDMP snapshot policy for wildcard backup fails with error 4201	278
	Troubleshooting with bpfis log	278
	Limitations of using HP-UX 11.31	278
Appendix A	Managing nbu_snap (Solaris)	281
	About managing nbu_snap	281
	Cache for nbu_snap	281
	About determining cache size	282
	About terminating nbu_snap	282
	nbu_snap commands	282
Appendix B	Overview of snapshot operations	287
	Introduction to snapshot operations	287
	Pre and post snapshot creation operations	288
	About quiescing the system	288
	About quiescing the database application	289
	About quiescing the stack	290
	File system quiesce	290
	Volume manager data caching	290
	How copy-on-write works	290
Index		295

Introduction

This chapter includes the following topics:

- Snapshot Client features at a glance
- Snapshot Client features
- 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

Snapshot Client features at a glance

NetBackup Snapshot Client provides a variety of snapshot-based features for NetBackup. It supports clients on UNIX, Linux, and Windows platforms, on Fibre Channel networks (SANs) or traditional LANs.

Table 1-1 Snapshot Client features at a glance

Snapshot Client feature	Description
Snapshot	<p>A point-in-time, read-only, disk-based copy of a client volume. NetBackup backs up data from the snapshot, not directly from the client's primary or original volume.</p> <p>Required by all features of Snapshot Client.</p>
Instant Recovery	<p>Makes the backups available for recovery from disk.</p>
Off-host backup	<p>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.</p> <p>Includes alternate the client, data mover, and virtual machine host (for VMware).</p> <p>Note: NetBackup 7.1 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.</p>
FlashBackup	<p>Combines the speed of raw-partition backups with the ability to restore individual files.</p>
NetBackup for Hyper-V	<p>Backs up and restores Windows and Linux Hyper-V virtual machines (guest operating systems).</p> <p>See the NetBackup for Hyper-V Administrator's Guide.</p>
NetBackup for VMware	<p>Backs up and restores Windows and Linux VMware virtual machines (guest operating systems).</p> <p>See the NetBackup for VMware Administrator's Guide.</p>
NAS snapshot	<p>Makes snapshot-based backup of data on a Network Attached Storage (NAS) host.</p>
Block level incremental backup (BLIB)	<p>Enables NetBackup to back up only the changed data blocks of VMware virtual machines and Oracle or DB2 database files.</p>

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

Snapshot Client feature	Description
NetBackup Replication Director	The implementation of NetBackup OpenStorage-managed snapshot replication, where the snapshots are stored on the storage systems of partnering companies. Replication is conducted per defined operations in storage lifecycle policies. See the NetBackup Replication Director Solutions Guide.

Snapshot Client features

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 snapshot methods

NetBackup can create different types of snapshots. Each snapshot type that you configure in NetBackup is called a snapshot method. Snapshot methods enable NetBackup to create snapshots within the storage stack (such as the file system, volume manager, or disk array) where the data resides. If the data resides in a logical volume, NetBackup can use a volume snapshot method to create the snapshot. If the data resides in a file system, NetBackup can use a file system method, depending on the client OS and the file system type.

Many different technologies are available for creating snapshots, and many different terms are used to refer to the underlying implementation of any given snapshot. Common terms include clone, split-mirror, and copy-on-write. In this documentation, the term "snapshot" designates any point-in-time, read-only copy of a primary

volume, regardless of its underlying implementation. Method-specific terminology is explained along with specific snapshot methods in other chapters of this guide.

In some cases, more than one method can make the snapshot. If the data resides in a file system over a logical volume, NetBackup could use a file system method or logical volume method. The choice of method might depend on the snapshot features available in the storage subsystem where the data resides. Or, the choice might depend on the requirements of the snapshot method itself. For example: if the client data is in a VxFS file system over a VxVM volume, NetBackup could create the snapshot with a file system method. On the other hand, NetBackup can use a volume manager method to create the snapshot of the same data, such as VxVM or FlashSnap. Between VxVM and FlashSnap, only FlashSnap supports the Persistent FastResync feature of VxVM mirror volumes. To take advantage of the Persistent FastResync feature, you would have to select the FlashSnap method.

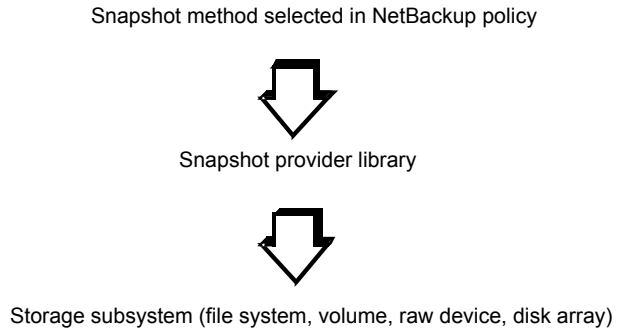
About snapshot providers

Each snapshot method relies on the snapshot technology that is built into the storage subsystem where the data is stored. Examples of storage subsystem are volume manager, file system, or hardware disk array. NetBackup includes a set of Software Libraries that are called "snapshot providers." The providers enable Snapshot Client to access the snapshot technology in the storage subsystem.

Each snapshot provider is designed for a particular subsystem. For example, the VxFS provider enables NetBackup to create snapshots of files in the Veritas File System (VxFS). The VxVM provider does the same for the data that is configured in Veritas Volume Manager volumes. The EMC CLARiiON disk array provider enables NetBackup to create hardware snapshots in the CLARiiON array.

You specify the method in the NetBackup policy. When the policy runs, the snapshot method calls the snapshot provider library. The provider then accesses the underlying commands in the storage subsystem to create the snapshot.

Figure 1-1 Simplified view of NetBackup access to snapshot technology

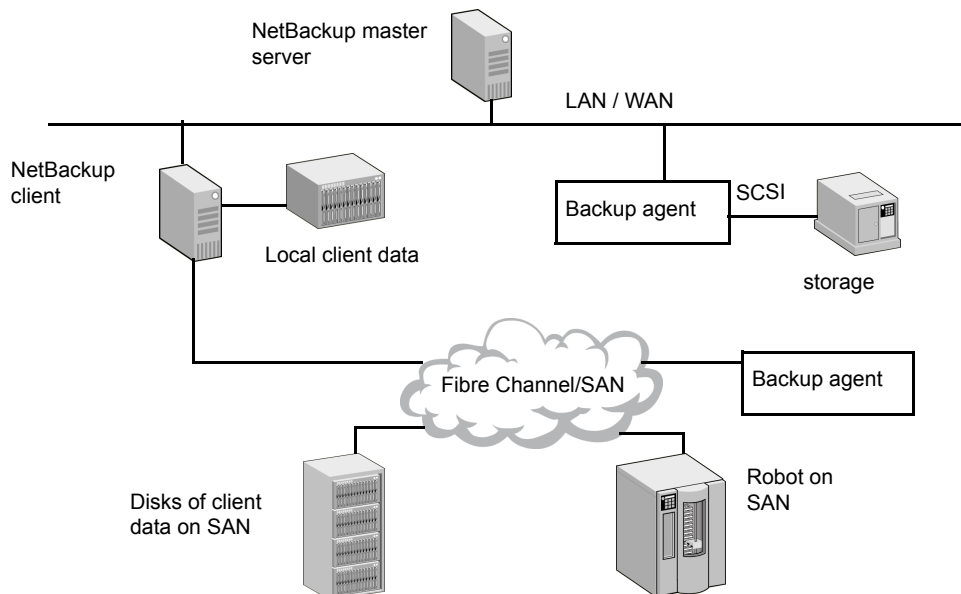


About off-host backup support

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 1-2 shows a backup agent.

Figure 1-2 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 7.1 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 234.

About FlashBackup policies

FlashBackup is a policy type that combines the speed of raw-partition backups with the ability to restore individual files.

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:

- 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 Client and NDMP

Using the NDMP protocol version V4 snapshot extension, NetBackup Snapshot Client can make policy-based snapshots of data on a Network Attached Storage (NAS) host. The snapshot is stored on the same NAS device that contains the primary client data. From the snapshot, you can restore individual files or roll back an entire volume or file system, by means of Instant Recovery.

Note: NetBackup for NDMP add-on software is required, and the NAS vendor must support snapshots.

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 25.

Two types of snapshots are available, both supported by NetBackup: copy-on-write and mirror (or clone).

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 transactionally consistent record.

See “How copy-on-write works” on page 290.

Note that VxFS allows two kinds of copy-on-write snapshots: file system snapshots and Storage Checkpoints.

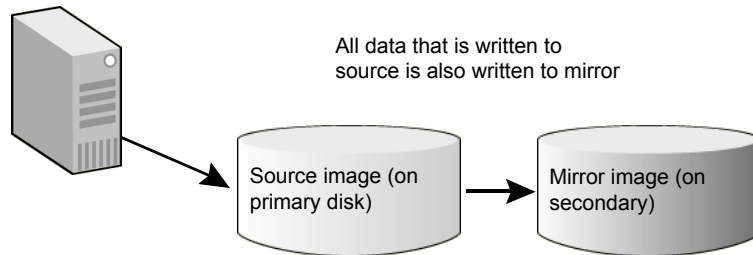
See “Benefits of copy-on-write versus mirror” on page 23.

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 1-3 Source is copied to mirror

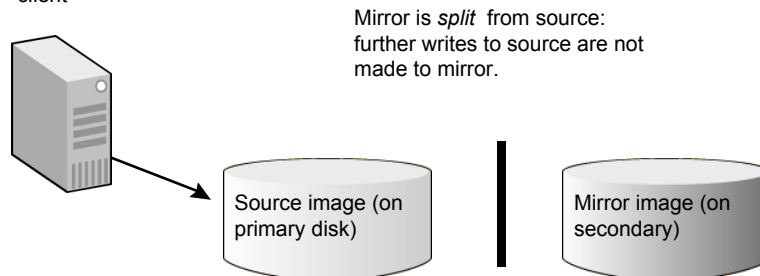
NetBackup
client



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 1-4 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 23.

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
<ul style="list-style-type: none">■ 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.	<ul style="list-style-type: none">■ 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. <p>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.</p>

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 1-5 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: Third-Party Copy Device data mover (UNIX clients only)	<p>A third-party copy device reads raw data from the client snapshot and writes the data to a storage device. To do so, the third-party copy device uses the Extended Copy command and mapping information from the client. Many kinds of devices, such as routers and disk arrays, are designed as third-party copy devices.</p> <p>A list of supported third-party copy devices is available. See “Snapshot Client assistance” on page 42.</p>
Data mover: NDMP	<p>Use to replicate NDMP snapshots. Select this agent in a policy that uses NDMP with Replication Director.</p> <p>For more information about configuring a policy to use NDMP with Replication Director, see the NetBackup Replication Director Solutions Guide.</p>
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

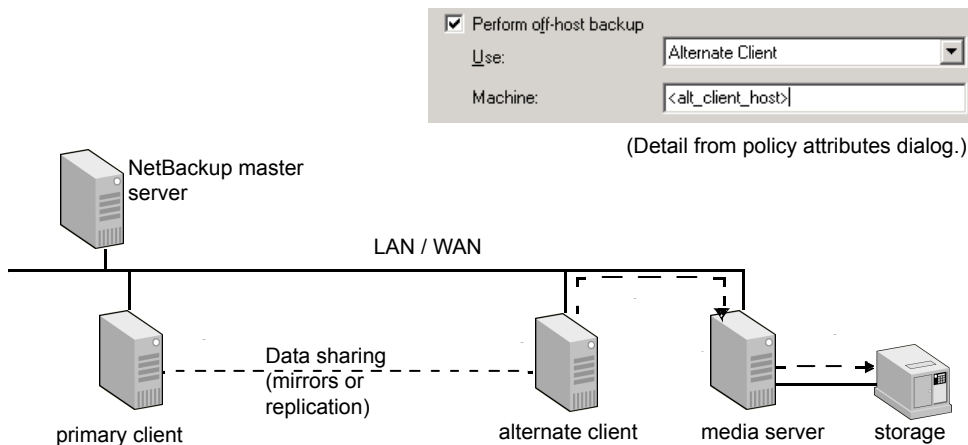
NetBackup Snapshot Client supports several forms of off-host backup, as explained in the following sections.

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 1-6 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

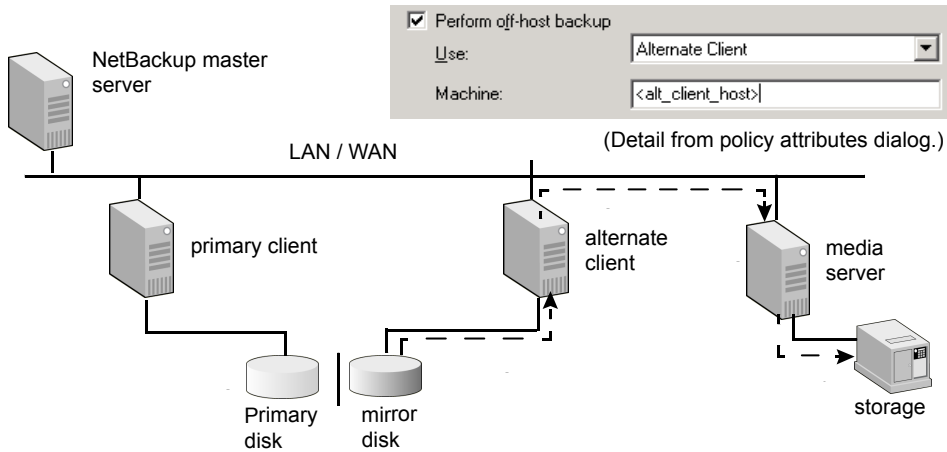
For alternate client backup, the original (primary) and alternate client must collaborate to create a snapshot. The following sections highlight two configurations: hardware array split-mirror snapshot and volume manager snapshot with data replication. Other configurations are possible, depending on 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 1-7 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 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 1-8 shows the media server and alternate client on the same host.


Alternate client and media server on same host

☒ Perform off-host backup

Use:

Machine:

(Detail from policy attributes dialog.)



The figure shows the following phases:

Phase	Action
Phase 1	Mirror disk is synchronized with primary.
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 1-9 Alternate client for multiple primary clients





Figure 1-10 Multiple clients with SSO: alternate client performs backup for multiple primary clients with NetBackup SSO option on a SAN

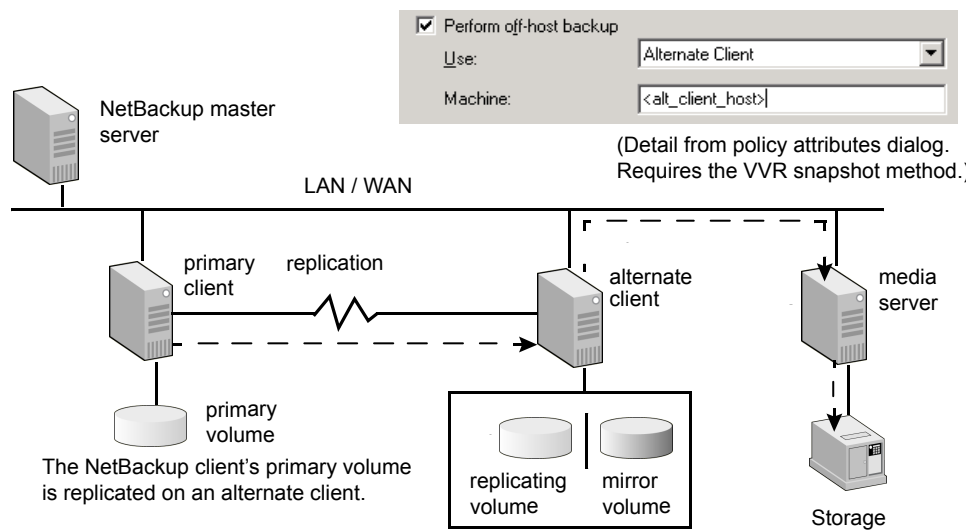


Alternate client backup through data replication example (UNIX only)

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 1-11 Replication: primary client and alternate client share data through replication



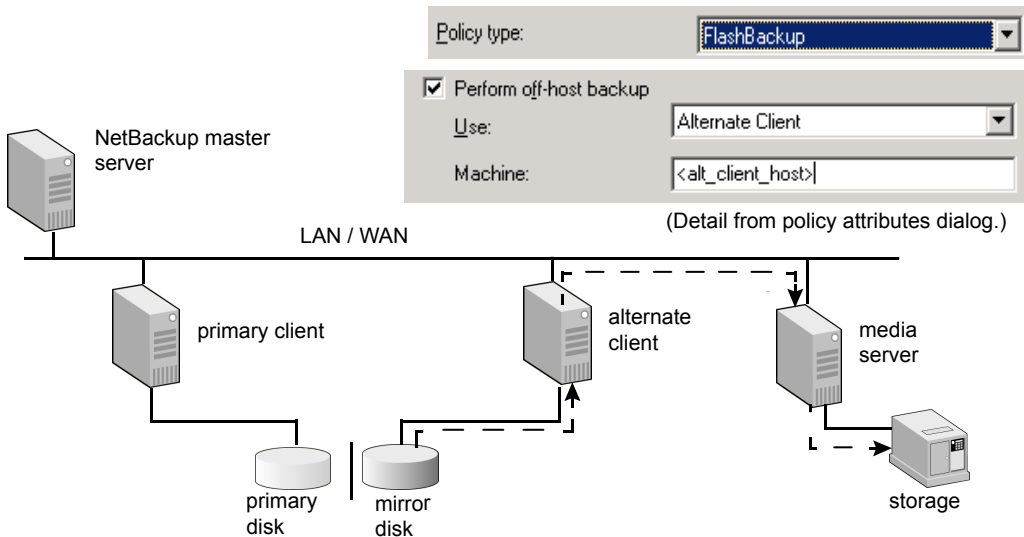
Phase	Action
Phase 1	Primary client collaborates with the alternate client to create snapshot.
Phase 2	Alternate client creates the snapshot by splitting the mirror from a replicating volume.
Phase 3	Alternate client sends the snapshot data from the snapshot to the media server.
Phase 4	Media server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

FlashBackup and alternate client combination example

FlashBackup is a policy type that combines the speed of raw-partition backup with the ability to restore individual files. FlashBackup can be combined with off-host backup methods, such as alternate client backup in the split-mirror configuration.

Figure 1-12 Alternate client split-mirror backup with FlashBackup policy type



The figure shows the following phases:

Phase	Action
Phase 1	Mirror disk is synchronized with primary.
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 creates file system map of the snapshot.
Phase 4	Alternate client sends the file system map to the media server.
Phase 5	Alternate client streams the snapshot-based backup from the mirror to the media server.

Phase	Action
Phase 6	Media server reads the backup image from the alternate client.
Phase 7	Media server writes the backup image to storage.

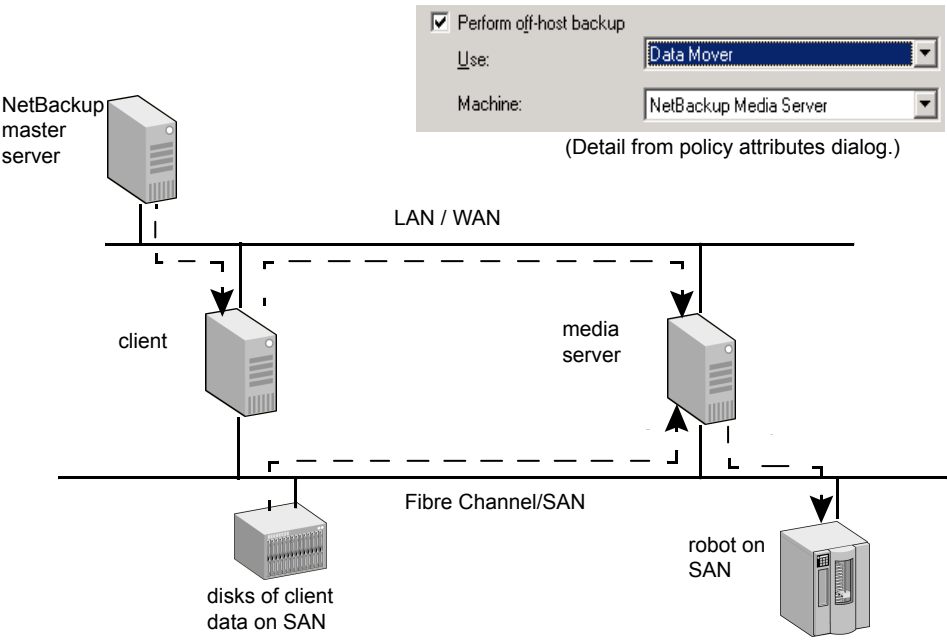
NetBackup media server data mover example (UNIX only)

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 1-13 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 4	Media server writes data across the SAN to storage.

About the NDMP data mover

Replication Director uses NDMP to perform the following operations:

- Back up snapshots
- Perform a live browse of snapshots.
- Restore from snapshots for the copy back method.

Support for all these operations is provided for replicated snapshots as well.

For more information about configuring a policy to that uses NDMP with Replication Director, see the NetBackup Replication Director Solutions Guide.

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 42.

Please note the following additional requirements:

- For the VxFS_Checkpoint snapshot method, all clients must have VxFS 3.4 or later with the Storage Checkpoints feature.

- To use Snapshot Client to back up a VxFS file system, the client's VxFS file system has to be patched with the dynamic linked libraries.
- For the VxVM snapshot method, all clients must have VxVM 3.1 or later.
- For the FlashSnap and VVR snapshot methods, all clients must have VxVM 3.2 or later. Each method requires its own add-on license to VxVM.
- 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.
- HP clients must use the OnlineJFS file system, not the default JFS.

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 42.

Note the following restrictions:

- Snapshot Client does not support the ALL_LOCAL_DRIVES entry in the policy's **Backup Selections** list.
- The VxFS_Checkpoint and VxVM snapshot methods support VxFS multi-volume file systems only.
- For the NetBackup media server or Third-Party Copy Device methods, the client disk must be either a SCSI or Fibre Channel device.
- For off-host backup that uses a data mover with the `nbu_snap`, `VxFS_Checkpoint`, or `VxVM` snapshot methods: The NetBackup media server must have access to all the disks that make up the snapshot. The disk(s) can be connected to a SAN. For each of these snapshot methods, note the following:
 - **nbu_snap**: Media server requires access to the active disk and the cache disk.
 - **VxFS_Checkpoint**: Media server requires access to the primary or an active disk.

- **VxVM:** Access requirements depend on layout of the volume group. Media server must be able to access all disks that make up the SnapMirror volume.
- Backup of an AIX 64-bit client with the NetBackup media server (data mover) method and the VxVM or VxFS_Checkpoint snapshot method may fail with NetBackup status code 11. This failure may occur if the client volumes are configured with Storage Foundation 5.0 MP3. A NetBackup message similar to the following appears in the job's Detailed Status tab:

```
12/09/2010 23:23:23 - Error bpbrm (pid=458874) from  
client p5201: ERR - bp_map_open, err 2059
```

This error occurs because the required VxVM libraries for 64-bit AIX are not installed in the correct location. The libraries should be installed in
`/opt/VRTSvxms/lib/map/aix64/.`

```
cp /usr/lpp/VRTSvxvm/VRTSvxvm/5.0.3.0/inst_root/  
opt/VRTSvxms/lib/map/aix64/* /opt/VRTSvxms/lib/map/aix64/
```

Note: This issue is fixed in Storage Foundation versions starting with 5.0MP3RP3, 5.1RP1, and 5.1SP1.

- For off-host backups that use the NDMP data mover option to replicate snapshots, see the NetBackup Replication Director Solutions Guide for a list of limitations.
- In a clustered environment, Instant Recovery point-in-time rollback is not supported for the backups that were made with a disk array snapshot method. The disk array snapshot methods are described in the chapter titled Configuration of snapshot methods for disk arrays. See “About the new disk array snapshot methods” on page 145.
- For the TimeFinder, ShadowImage, or BusinessCopy legacy snapshot methods (when you use the NetBackup media server or Third-Party Copy Device backup methods): The NetBackup clients must have access to the mirror (secondary) disk containing the snapshot of the client's data. The NetBackup clients must also be able to access the primary disk. The NetBackup media server only needs access to the mirror (secondary) disk.
- For the TimeFinder, ShadowImage, or BusinessCopy legacy snapshot methods, a Volume Manager disk group must consist of disks from the same vendor.
- 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.

- For the NetBackup media server or Third-Party Copy Device backup method: The disk must return its SCSI serial number in response to a serial-number inquiry (serialization), or the disk must support SCSI Inquiry Page Code 83.
- Multiplexing is not supported for Third-Party Copy Device off-host backups.
- For alternate client backup: The user and the group identification numbers (UIDs and GIDs) for the files must be available to the primary client and the alternate backup client.
- Inline Tape Copies (called Multiple Copies in Vault) is not supported for Third-Party Copy Device off-host backups.
- For media servers running AIX (4.3.3 and higher), note the following:
 - Clients must be Solaris, HP, or AIX.
 - Requires the use of tape or disk LUNs to send the `Extended copy` commands for backup.
 - The tape must be behind a third-party-copy-capable FC-to-SCSI router. The router must be able to intercept Extended Copy commands that are sent to the tape LUNs.
 - The `mover.conf` file must have a tape path defined, not a controller path.

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
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.

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

Term	Definition
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 290.
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.
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	<p>A contiguous set of disk blocks that are allocated for a file and represented by three values:</p> <ul style="list-style-type: none"> ■ Device identifier ■ Starting block address (offset in the device) ■ Length (number of contiguous blocks) <p>The mapping methods in Snapshot Client determine the list of extents and send the list to the backup agent.</p>
FastResync (VxVM)	Formerly known as Fast Mirror Resynchronization or FMR, VxVM FastResync performs quick and efficient resynchronization of mirrors. NetBackup's Instant Recovery feature uses FastResync to create and maintain a point-in-time copy of a production volume.
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.

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

Term	Definition
File system	Has two meanings. For a product, such as UFS (Sun Solaris) or VxFS (Veritas) file systems, file system means the management and the allocation schemes of the file tree. Regarding a file tree component, file system means a directory that is attached to the UNIX file tree by means of the <code>mount</code> command. When a file system is selected as an entry in the NetBackup Backup Selections list, this definition applies.
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. <ul style="list-style-type: none"> ■ 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.
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.

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

Term	Definition
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.
Snapshot	A point-in-time, read-only, disk-based copy of a client volume. A snapshot is created with minimal effect on other applications. NetBackup provides several types, depending on the device where the snapshot occurs: copy-on-write, mirror, clone, and snap.
Snapshot method	A set of routines for creating a snapshot. You can select the method, or let NetBackup select it when the backup is started (auto method).
Snapshot mirror	A disk mirror created by the Veritas Volume Manager (VxVM). Snapshot mirror is an exact copy of a primary volume at a particular moment, reproduced on a physically separate device.
Snapshot source	The entity (file system, raw partition, or logical volume) to which a snapshot method is applied. NetBackup automatically selects the snapshot source according to the entries in the policy's Backup Selections list.
Snapshot volume	A mirror that has been split from the primary volume or device and made available to users. Veritas Volume Manager (VxVM) creates snapshot volumes as a point-in-time copy of the primary volume. Subsequent changes in the primary volume are recorded in the Data Change Log. The recorded changes can be used to resynchronize with the primary volume by means of VxVM FastResync. The changes that were made while the snapshot volume was split are applied to the snapshot volume to make it identical to the primary volume.
Standard device	Refers to the primary disk in an EMC primary-mirror disk array (see primary disk).
Storage Checkpoint (VxFS)	<p>Provides a consistent and a stable view of a file system image and keeps track of modified data blocks since the last checkpoint. Unlike a mirror, a VxFS Storage Checkpoint does not create a separate copy of the primary or the original data. It creates a block-by-block account that describes which blocks in the original data have changed from the instant the checkpoint was activated.</p> <p>A Storage Checkpoint stores its information in available space on the primary file system, not on a separate or a designated device. (Also, the <code>ls</code> command does not list Storage Checkpoint disk usage; you must use the <code>fsckptadm list</code> command instead.)</p>

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

Term	Definition
Third-Party Copy Device	<p>Has two meanings:</p> <ul style="list-style-type: none">■ 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.
UFS file system	The UNIX file system (UFS), which is the default file system type on Sun Solaris. The UFS file system was formerly the Berkeley Fast File System.
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.
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.
VxFS	The Veritas extent-based File System (VxFS), designed for high performance and large volumes of data.
VxVM	The Veritas Volume Manager (VxVM), which provides the logical volume management that can also be used in SAN environments.

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 .

Table 1-5 Types of assistance available for Snapshot Client (*continued*)

Type of assistance	Description
Snapshot Client assistance from the web	<p>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:</p> <p>http://www.veritas.com/docs/000081320</p> <p>This document includes the following:</p> <ul style="list-style-type: none"> ■ 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	<p>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:</p> <p>http://www.netbackup.com/compatibility</p>
NDMP information on the web	<p>The Veritas Support website has a pdf document on supported NDMP operating systems and NAS vendors. The document also contains configuration and troubleshooting help for particular NAS systems.</p> <p>This document may be accessed from the following link:</p> <p>http://www.veritas.com/docs/000027113</p> <p>The document's title is: <i>NetBackup for NDMP Supported OS and NAS Appliance Information</i>.</p>

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
- Adding a Snapshot Client license key on UNIX
- Adding a Snapshot Client license key on Windows
- 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 8.0 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 or Windows platform. For a list of supported platforms for Snapshot Client, refer to the following document:
<http://www.veritas.com/docs/000081320>
- For a list of supported platforms for NetBackup for VMware and NetBackup for Hyper-V, refer to the following document:
<http://www.veritas.com/docs/000035749>
- NetBackup 8.0 or later client software must be installed on clients. For AIX and Linux clients, the client data must be in a VxFS file system.

- For Instant Recovery using the VxFS_Checkpoint method, the VxFS file system with the Storage Checkpoints feature must be installed on clients.

Snapshot Client installation notes

The following list contains Snapshot Client installation information for UNIX and Windows.

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 8.0, 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.

Windows

- For Windows, NetBackup Snapshot Client software is automatically installed with the core NetBackup server and client product.
For the NetBackup installation procedure, see the NetBackup Installation Guide.
- If you install Snapshot Client 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.

Note: Before 7.0, it was possible to uninstall Snapshot Client. From 7.0 however, there is no separate uninstall procedure for Snapshot Client as the software is automatically installed with the core NetBackup server. If you want to uninstall Snapshot Client, you must uninstall NetBackup completely.

Refer to the uninstall procedure that is described in the NetBackup Installation Guide

Adding a Snapshot Client license key on UNIX

To install Snapshot Client, you must add a valid license key, as follows.

To add a Snapshot Client license key on UNIX

- 1 Log on as root on the NetBackup master server.
- 2 Enter the following command to list and add keys:

```
/usr/opensv/netbackup/bin/admincmd/get_license_key
```

- 3 In a clustered environment, the previous steps must be done on each node in the cluster.
- 4 In a cluster environment, unfreeze the active node.

See the clustering section in the NetBackup High Availability Administrator's Guide for the cluster software you are running.

Adding a Snapshot Client license key on Windows

You must add a valid license key for Snapshot Client on each master server, as follows.

To add a Snapshot Client license key on Windows

- 1 Log on.
- 2 In the NetBackup Administration Console, choose **Help > License Keys**.
The **NetBackup License Keys** window appears. Existing keys are listed in the lower part of the window.
- 3 Click the star icon to open the **Add a New License Key** dialog.
- 4 Type the new license key in the **New license key** field and click **Add**.
The new license key appears in the lower part of the dialog box.
- 5 In a clustered environment, the previous steps must be done on each node in the cluster.
- 6 In a cluster environment, unfreeze the active node.

See the clustering section in the NetBackup High Availability Administrator's Guide for the cluster software you are running.

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. See “Adding a Snapshot Client license key on UNIX” on page 47.
- If you have a UNIX server with Windows clients, you must install the client software on the Windows client computers individually, from the NetBackup media. See “Adding a Snapshot Client license key on Windows” on page 47.

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 254.

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 Master server. In a cluster, if the client experiences problems and failover to another node occurs, the state file on the NetBackup Master server comes in handy to obtain information about the failed client. In case the state file does not exist on the NetBackup Master server, it becomes impossible for the active node to get snapshot information.

Table 2-1 Location of snapshot state file

Location	UNIX	Windows
On the client	/usr/opensv/netbackup/online_util/fi_cntl/	install_path\NetBackup\online_util\fi_cntl\
On the master server	/usr/opensv/netbackup/db/snapshot/client_name/	install_path\NetBackup\db\snapshot\client_name\

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
- Automatic snapshot selection
- 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 Enterprise and the Snapshot Client add-on product must be installed on master server(s) and clients.
- For the NetBackup Media Server and Third-Party Copy Device 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.
- FlashBackup policies do not support encryption or compression.
- BLIB with Snapshot Client (**Perform block level incremental backups** option on the policy **Attributes** tab): BLIB is supported with NetBackup for Oracle, NetBackup for DB2, and with VMware.
If you choose the **Perform block level incremental backups** option on the policy **Attributes** tab, the other features of Snapshot Client are grayed out.

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 master 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:

Data classification: <No data classification>

Policy storage: Any_available

Policy volume pool: NetBackup

Take checkpoints every: 0 minutes

Limit jobs per policy: 1

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:
 - For VMware virtual machines, select **FlashBackup-Windows VMware** or **MS-Windows**. For Hyper-V, select **FlashBackup-Windows** or **Hyper-V**. Refer to the appropriate NetBackup guide for your virtual environment.

- If client data is in a database, select the database type (**DB2, Oracle, MS-Exchange-Server, MS-SQL-Server, or SAP**).
NetBackup 8.0 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 System Administrator's Guide*: <https://www.veritas.com/docs/DOC5332>
- To use FlashBackup, select **FlashBackup** for UNIX clients or **FlashBackup-Windows** for Windows clients.

Note: 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).

- For all other cases, select **Standard** for UNIX clients and **MS-Windows** for Windows clients.
- 4 Select a storage unit, storage unit group, or a storage lifecycle policy as the **Policy storage**.
 - 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 Optional: select the snapshot method manually.
See "Selecting the snapshot method" on page 60.
Skip this step if you want NetBackup to select the snapshot method.
See "Automatic snapshot selection" on page 59.

- 7** To create a backup that enables Instant Recovery, 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 234.
Help for creating a policy for instant recovery backups is available.
See “Configuring a policy for Instant Recovery” on page 102.
- 8** To reduce the processing load on the client, select **Perform off-host backup**.
See “Off-host backup configuration options ” on page 57.
- 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 55.
- 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 76.
The NetBackup Administrator’s Guide, Volume I, describes other file-path rules.
- When you configure a snapshot method in a **MS-Windows** policy, the backward slash (\) must be entered in the **Backup Selections** list after the drive letter.

If the backward slash is not included, the snapshot image does not appear in the NetBackup catalog.

- 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 76.

- For a raw partition backup of a UNIX client, specify the `/rdsk` path, not the `/dsk` path. You can specify the disk partition (except on AIX) or a VxVM volume. Examples:

On Solaris: `/dev/rdsk/c0t0d0s1`
`/dev/vx/rdsk/volgrp1/vol1`

On HP: `/dev/rdsk/c1t0d0`
`/dev/vx/rdsk/volgrp1/vol1`

On AIX and Linux: `/dev/vx/rdsk/volgrp1/vol1`

On Linux: `/dev/sdc1`

On AIX clients, backing up a native disk partition is not supported. A raw partition backup must specify a VxVM volume, such as `/dev/vx/rdsk/volgrp1/vol1`.

Note that `/dev/vx/dsk/volgrp1/vol1` (without the “r” in `/rdsk`) does not work.

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 52.

The screenshot shows a dialog box titled "Snapshot Client and Replication Director". It contains several configuration options:

- ☐ Perform block level incremental backups
- ☐ Use Replication Director
- ☒ Perform snapshot backups (with an "Options..." button to its right)
- ☐ Retain snapshot for Instant Recovery or SLP management
- ☐ Hyper-V server: (with an empty text box to its right)
- ☒ Perform off-host backup

Below the "Perform off-host backup" option, there are two fields:

- Use:** A pull-down menu currently showing "Alternate client".
- Machine:** A pull-down menu currently showing "VMware backup host".

The "Machine" pull-down menu is open, showing three options: "Alternate client", "Data mover", and "VMware backup host".

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

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

- **VMware backup host**

Select this option to designate a VMware backup host as the backup agent, for VMware virtual machine (guest OS) backup. During restore, this host is called the VMware restore host.

Select the name of the backup host in the **Machine** pull-down.

To back up a virtual machine that does not have a NetBackup client installed on it, you must select this option. If a NetBackup client is installed on the virtual machine, you can back up the virtual machine in the same way as an ordinary physical host (a snapshot-based backup is not required).

The **VMware backup host** option requires the FlashBackup-Windows or MS-Windows policy type.

See the *NetBackup for VMware Administrator's Guide* for further information:
<https://www.veritas.com/docs/DOC5332>

Note: The VMware backup host is not displayed when you select the **Retain snapshots for Instant Recovery or SLP management** check box as VMware backup is not supported for Instant Recovery.

■ **Alternate Client**

Select this option to designate another client (alternate client) as the backup agent.

An alternate client saves computing resources on the original client. The alternate client handles the backup I/O processing on behalf of the original client, so the backup has little effect on the original client.

Enter the name of the alternate client in the **Machine** field.

See “About using alternate client backup” on page 73.

■ **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** pull-down:

Network Attached Storage An NDMP host (NAS filer) performs the backup processing, by means of the NAS_Snapshot method. NetBackup for NDMP software is required on the NetBackup server. This option is required for NAS snapshots.

See “Setting up a policy for NAS snapshots” on page 120.

NetBackup Media Server A Solaris, HP, AIX media server performs the backup processing (for Solaris, HP, and AIX clients only).

Third-Party Copy Device

A third-party copy device handles the backup processing. For Solaris, HP, AIX, and Linux clients only.

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.

For more information on media server and third-party copy backup, refer to the NetBackup Snapshot Client Configuration document:

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

The use of **Perform off-host backup** may require additional configuration, as follows:

- For the policy storage unit or group of storage units, note:
Any_available is not supported for the following data mover types: NetBackup Media Server and Third-Party Copy Device. Disk storage units are not supported for Third-Party Copy Device.
Instead of using a particular storage unit, you can create a storage unit group that designates devices that are configured on the SAN.
See the *NetBackup Administrator's Guide, Volume I*:
<https://www.veritas.com/docs/DOC5332>
- Also note the following:
 - For the Third-Party Copy Device option, you must create a `mover.conf` file. Assistance with this file is available, in the NetBackup Snapshot Client Configuration document:
<http://www.veritas.com/docs/000081320>
 - If you do not have Veritas CommandCentral Storage and your backup devices do not support identification descriptors (SCSI E4 target descriptors), you may need to edit the `3pc.conf` file.
See <http://www.veritas.com/docs/000081320>

Automatic snapshot selection

To have NetBackup select the snapshot method, click **Perform snapshot backups** on the policy **Attributes** tab. Note the following points:

- If the policy is new, NetBackup selects a snapshot method when the backup starts (by default, the snapshot method is set to **auto**).

- If the policy had been configured for a particular snapshot method, click the **Snapshot Client Options** option and set the snapshot method to that particular one. NetBackup selects a snapshot method when the backup starts.

Use of the auto method does not guarantee that NetBackup can select a snapshot method for the backup. NetBackup looks for a suitable method according to the following factors:

- The client platform and policy type.
- The presence of up-to-date software licenses, such as VxFS and VxVM.
- How the client data is configured. For instance:
 - Whether a raw partition has been specified for a copy-on-write cache. See “Entering the cache” on page 127.
 - Whether the client’s data is contained in the VxVM volumes that were configured with one or more snapshot mirrors.

NetBackup uses the first suitable method found.

Note: The auto method cannot select a snapshot method that is designed for a particular disk array, such as EMC_TimeFinder_Clone or HP_EVA_Vsnap. You must select the disk array method from the drop-down list on the **Snapshot Options** dialog box.

The auto method has the following advantages:

- NetBackup can use a different snapshot method for each item in the **Backup Selections** list, or for each client in the policy. As a result, NetBackup has more flexibility in choosing a snapshot method, and you can circumvent the one-snapshot-method-per-policy restriction.
- At the time of backup, NetBackup selects a snapshot method according to how the client is configured at that moment. If the VxFS or VxVM license on the client has expired, or if the client was recently reconfigured, NetBackup selects a snapshot method accordingly.

Selecting the snapshot method

You can manually select a snapshot method as follows.

To select a snapshot method

- 1 On the master 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**.

Snapshot Options - Policy test5

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 the **Snapshot method** for the policy.
 - Choose **auto** if you want NetBackup to select the snapshot method.
See “Automatic snapshot selection” on page 59.
 - The available methods depend on how your clients are configured and which attributes you selected on the **Attributes** tab.

Only one snapshot method can be configured per policy. Configure each policy for a single method and include only clients and backup selections for which that snapshot method can be used. For example, for the nbu_snap method (which applies to Solaris clients only), create a policy that includes Solaris clients only. The snapshot method you select must be compatible with all items in the policy’s **Backup Selections** list.

See “Snapshot methods” on page 62.

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

 See “Configuration parameters for Snapshot Client” on page 65.

 See “Snapshot Resources” on page 71.
- 8 Click **OK**.

Snapshot methods

Table 3-1 describes each snapshot method (not including the disk array methods).
See “Disk array methods at a glance” on page 149.

Table 3-1 Snapshot method descriptions

Method	Description
auto	NetBackup selects a snapshot method when the backup starts. If necessary, NetBackup selects a different method for each item in the Backup Selections list.
FlashSnap	For mirror snapshots on alternate clients, with the VxVM FlashSnap feature. This method is for clients on Solaris, HP, AIX, Linux, and Windows. UNIX clients must be at VxVM 3.2 or later. Linux and AIX clients must be at VxVM 4.0 or later. Windows clients must be at VxVM 3.1 or later, with all the latest VxVM service packs and updates. FlashSnap is based on the VxVM disk group split and join technology.
Hyper-V	For snapshots of Hyper-V virtual machines. Refer to the NetBackup for Hyper-V Administrator's Guide: https://www.veritas.com/support/en_US/article.DOC5332
NAS_Snapshot	For copy-on-write snapshots of client data that resides on an NDMP host. Requires NetBackup for NDMP software. Further information is available on requirements and configuration details. See “About NAS snapshot overview” on page 117.

Table 3-1 Snapshot method descriptions (*continued*)

Method	Description
nbu_snap	<p>For copy-on-write snapshots of UFS or Veritas VxFS file systems. For Solaris clients only. nbu_snap is not supported in clustered file systems.</p> <p>NetBackup does not currently support extended volume table of contents (VTOCs) or EFI labels when nbu_snap is used with the UNIX file system (UFS). NetBackup supports an effective limit of 1TB disks (regardless of partition size) when UFS and nbu_snap are used together. To use nbu_snap with a larger disk (greater than 1TB), you must use VxVM instead of UFS.</p> <p>This method requires a raw partition that is designated for cache.</p>
OST_FIM	<p>For use in a policy that is configured for snapshot replication using Replication Director. The OST_FIM method works only if the Policy storage attribute points to a storage lifecycle policy that contains a snapshot-capable storage unit.</p>
VMware	<p>For snapshots of VMware virtual machines.</p> <p>Refer to the <i>NetBackup for VMware Administrator's guide</i></p>
VSP	<p>VSP is currently deprecated. VSP is for snapshots of open and active files on NetBackup pre-7.0 Windows clients. For 7.0 and later versions, clients in a policy that is configured with VSP, NetBackup automatically uses VSS instead of VSP.</p> <p>For more information on this method (Veritas Volume Snapshot Provider), refer to the NetBackup Administrator's Guide, Volume I.</p> <p>You can use VSP without Snapshot Client. In some cases, however, such as when the Busy File Timeout has expired, no snapshot is created. The backup job then may continue without backing up the busy file. If you use VSP with Snapshot Client, the backup successfully creates a snapshot of all files or the backup job fails.</p>

Table 3-1 Snapshot method descriptions (*continued*)

Method	Description
VSS	<p>VSS uses the Volume Shadow Copy Service of Windows and supports Instant Recovery. VSS is for local backup or alternate client backup.</p> <p>For the most up-to-date list of Windows operating systems and disk arrays supported by this method, see the following NetBackup Snapshot Client Configuration document: https://www.veritas.com/support/en_US/article.000081320</p> <p>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.</p> <p>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.</p>
VVR	<p>For alternate client backups of a replicated VxVM volume. For clients on Solaris, HP, Linux, AIX.</p> <p>Requires VxVM 3.2 or later with the Veritas Volume Replicator license. Linux and AIX clients require VxVM 4.0 or later.</p>
VxFS_Checkpoint	<p>For copy-on-write snapshots of clients on Solaris, HP, AIX, or Linux. The FlashBackup policy type does not support this method.</p> <p>Requires the Storage Checkpoint feature of VxFS 3.4 or later. HP requires VxFS 3.5 or later. Linux and AIX clients require VxFS 4.0 or later.</p> <p>Note that VxFS_Checkpoint requires the Snapshot Client license and the Veritas File System license with the Storage Checkpoints feature. Without both licenses, the copy-on-write snapshot (Storage Checkpoint) cannot be opened and the backup fails.</p>

Table 3-1 Snapshot method descriptions (*continued*)

Method	Description
VxFS_Snapshot	<p>For copy-on-write snapshots of Solaris or HP clients on the local network (not off-host), for FlashBackup policies only. This method requires VxFS 3.4 (Solaris) or VxFS 3.3.2 (HP) or later. This method also requires a designated cache.</p> <p>Note that all files in the Backup Selections list must reside in the same file system.</p>
VxVM	<p>For any of the following types of snapshots with data configured over Volume Manager volumes, for clients on Solaris, HP, AIX, Linux, or Windows. Linux and AIX clients require VxVM 4.0 or later.</p> <p>For "third-mirror" snapshots (VxVM 3.1 or later).</p> <p>For full-sized instant snapshots (VxVM 4.0).</p> <p>For space-optimized instant snapshots (VxVM 4.0).</p>
remote_vxfs	<p>Creates a vxfs_checkpoint snapshot (copy-on-write) of database backup shares on the NetBackup appliance.</p> <ul style="list-style-type: none">■ This method is only applicable for Oracle Agent.■ Restore from snapshot is only done through copy back method. <p>For more information about remote_vxfs, see the NetBackup Administrator's Guide for Oracle.</p>

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.

Cache device path parameter

Specify a raw partition for the cache (logical volume or a physical disk) by entering the cache partition's full path name in the **Value** field. For example:

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

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

HP LVM volume: `/dev/volume_group_1/volume_3`

This setting overrides a cache that is specified on **Host Properties > Clients > Client Properties dialog > UNIX Client > Client Settings**.

See “Entering the cache” on page 127.

Do not specify wildcards (such as `/dev/rdisk/c2*`).

A complete list of requirements is available.

See “Cache device requirements” on page 124.

Warning: The cache partition’s contents are overwritten by the **nbu_snap** or **VxFS_Snapshot** process.

Delay in seconds between disk group split retries parameter

Certain routine operating system processing must complete before the snapshot volume is available and the disk group can be split. By default, NetBackup waits 60 seconds before it retries the disk group split if **Number of times to retry disk group split** is 1 or more. On some systems, a 60-second delay may be too short. Use this parameter to set a longer delay between retries.

IBC receive timeout parameter (seconds)

Determines how long NetBackup waits for the next end-of-transaction marker to be received in the VVR replication data stream. For instance, a process may fail to send the marker, or the network connection may be down or overloaded. If this timeout is exceeded before the next marker is received, the backup fails.

IBC send timeout parameter (seconds)

Determines how long NetBackup waits for the next end-of-transaction marker to be automatically removed from the VVR replication data stream. If the marker is not removed, replication cannot continue. If this timeout is exceeded before the current marker is removed, the backup fails.

Keep snapshot after backup parameter

This option specifies whether or not the snapshot image is retained on the mirror disk after the backup completes (default is **No**). Retaining the image (**Yes**) enables a quick restore from the mirror in case data is deleted from the client’s primary disk.

The image is retained on the mirror until the next time a backup is run using this policy. **No** indicates that the image is deleted from the mirror as soon as the backup completes.

If the client is restarted, the snapshots that have been kept must be remounted before they can be used for a restore. You can use the `bpfis` command to discover the images.

See the `bpfis` man page or the *NetBackup Commands* manual.

Note however that NetBackup automatically remounts Instant Recovery snapshots.

If the snapshot is on an EMC, Hitachi, or HP disk array, and you want to use hardware-level restore, important information is available.

See the Warning under Hardware-level disk restore in the *NetBackup Snapshot Client Configuration* document. This document may be accessed from the following location:

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

Maximum number of volumes to resynchronize concurrently parameter

For the configurations that have sufficient I/O bandwidth, multiple volumes can be resynchronized simultaneously, to complete resynchronization sooner. This parameter specifies the number of volume pairs that are resynchronized simultaneously. A volume pair consists of a source volume and its snapshot (or mirror) volume.

The default is 1, which means that volume pairs are resynchronized one at a time. Accept the default if the I/O bandwidth in your clients and disk storage cannot support simultaneous synchronization of volumes. A major factor in I/O bandwidth is the number and speed of HBAs on each client.

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 master server, it is advisable that you validate and correct the policy after you upgrade to a NetBackup 8.0 master server.

See "Means of controlling snapshots" on page 100.

Number of times to retry disk group split parameter

After the FlashSnap snapshot is created, the disk group must be split to make the snapshot data accessible to the alternate client. For a short time, certain routine operating system processing must complete before the snapshot volume is available and the disk group can be split. The disk group cannot be split before the OS has made the volume available.

Use this parameter to set the number of times to retry the disk group split. The default is 0 (no retries): if the first attempt to split the disk group does not succeed, the backup fails.

You can also set the **Delay in seconds between disk group split retries** option.

Provider Type parameter

Indicates the type of VSS snapshot provider that creates the snapshot.

0-auto (the default)

If the policy is not configured for Instant Recovery, you can select this option. The auto option attempts to select the available provider in this order: Hardware, Software, System.

For an Instant Recovery backup, you must select the appropriate provider type (not **0-auto**) and snapshot attribute; otherwise, the backup fails.

For example: To use a hardware provider to back up a CLARiiON array using the EMC_CLARiiON_SnapView_Clone method, you must select **3-hardware** as the provider type and **2-plex** as the snapshot attribute.

1-system

Use the Microsoft system provider, for a block-level copy on write snapshot.

- Unlike options 2 and 3, the Microsoft system provider does not require any additional software or hardware.
- The Microsoft system provider does not support off-host backup, such as alternate client backup. For off-host backup, select option 2 or 3, depending on your environment.

2-software

For instance, use the Veritas Storage Foundations provider, for VxVM volumes. The software provider intercepts I/O requests at the software level between the file system and the volume manager.

3-hardware

Use the hardware provider for your disk array. A hardware provider manages the VSS snapshot at the hardware level by working with a hardware storage adapter or controller. For example, if you want to back up an EMC CLARiiON or HP EVA array by means of the array's snapshot provider, select **3-hardware**.

- Depending on your array and on the snapshot attribute you select, certain preconfiguration of the array may be required.
See the appropriate topic for your disk array and snapshot method.
See "Important note on VSS and disk arrays" on page 71.

Resynchronize mirror in background parameter

This option determines whether or not the backup job completes before the resynchronize operation has finished (default is **No**). **Yes** means that a resynchronize request is issued, but the backup does not wait for the resync operation to complete. **No** means that the backup job cannot reach completion until the resynchronize operation has finished.

Choosing **Yes** may allow more efficient use of backup resources. If two backups need the same tape drive, the second can start even though the resynchronize operation for the first job has not completed.

Snapshot Attribute parameter

Indicates the type of VSS snapshot to be created.

0-unspecified

If the policy is not configured for Instant Recovery, you can select this option.

For an Instant Recovery backup, you must select a snapshot attribute of 1 or 2; otherwise, the backup fails.

1-differential

Use a copy-on-write type of snapshot. For example, to back up an EMC CLARiiON array using the EMC_CLARiiON_SnapView_Snapshot method, select differential.

2-plex

Use a clone or a mirror type of snapshot. For example, to back up an EMC CLARiiON array using the EMC_CLARiiON_SnapView_Clone method, select plex.

Sync I/O delay in milliseconds parameter (UNIX only)

Specifies the delay in milliseconds between synchronizing successive sets of regions as specified by the value of the previous parameter (Sync region size in MB). You can use this option to change the effect of synchronization on system performance.

The default is 0 milliseconds (no delay). Increasing this value slows down synchronization. It also reduces the competition for I/O bandwidth with other processes that may access the volume at the same time.

This option is the same as the `slow=iodelay` parameter on the VxVM `vxsnap` command. For more details on the `slow=iodelay` parameter, see the *Veritas Volume Manager Administrator's Guide*.

Sync region size in MB parameter (UNIX only)

Specifies the size in MB of each I/O request that is used when the regions of a volume are synchronized. Specifying a larger size causes synchronization to complete sooner, but with greater effect on the performance of other processes that access the volume. 1 MB (the default) is the minimum value that is suggested for high-performance array and controller hardware. The specified value is rounded to a multiple of the volume's region size.

This option is the same as the `iosize=size` parameter on the VxVM `vxsnap` command. For more details on the `iosize=size` parameter, see the *Veritas Volume Manager Administrator's Guide*.

Wait for mirror sync completion parameter

This parameter applies only to the off-host backups that use the NetBackup Media Server or Third-Party Copy Device backup method.

With the default setting 1 (Yes), this parameter specifies that full-sized instant snapshots are not available for backup until the mirror synchronization is complete. Before synchronization is complete, some of the data that is required for the backup resides on the source disks but not on the snapshot disks. If the media server has access to the snapshot disks but does not have access to the source disks, and if the backup starts before the snapshot disks are fully synchronized with the source, the backup fails.

For a NetBackup Media Server or Third-Party Copy Device backup—if the media server cannot access the source disks—set this parameter to 1. If both the source and the snapshot disks are visible to the media server, you can set this parameter to 0 (No).

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.

Synchronize mirror before the backup parameter

Determines whether or not the primary and the mirror devices are automatically synchronized (if they were not already synchronized) before the backup begins. Default is **No**.

Specify **Yes** to have unsynchronized devices synchronized before the backup begins. **No** means that unsynchronized devices are not synchronized before the backup starts. In this case (**No**), if the primary and the mirror devices are not synchronized, the backup cannot succeed.

Snapshot Resources

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

See “Disk array configuration tasks” on page 151.

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 master 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 (such as clones, mirrors, or replication disks).
- For the FlashSnap and VVR snapshot methods, the following must be installed: VxVM 3.2 or later for UNIX, VxVM 4.0 or later for Linux and AIX, or VxVM 3.1 or later for Windows. Also, volumes must be configured over the primary host's disks. The VxVM FlashSnap or VVR license must also be installed.
- 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.
- Alternate client backup on Windows does not support incremental the backups that are based on archive bit. Instead, use incremental backups that are based on timestamp.
See "About incremental backup of mirror-based snapshots" on page 78. for more information.
- 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 3-2 lists the supported configurations.

Table 3-2 Alternate client requirements

If primary client is:	Alternate client must be:
Windows	Windows, at same level as primary client or higher
Solaris	Solaris, at same level as primary client or higher

Table 3-2 Alternate client requirements (*continued*)

If primary client is:	Alternate client must be:
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
VxFS 3.4 or later (VxFS 3.3 for HP, VxFS 4.0 for AIX and Linux)	VxFS, at same level as primary client or higher
VxVM 3.2 or later (UNIX) VxVM 3.1 or later (Windows)	<p>VxVM, at same level as primary client or higher.</p> <p>Note: For the VVR method, the alternate client must be at exactly the same level as primary client.</p> <p>For VxVM on Windows, use VxVM 3.1 or later with all the latest VxVM service packs and updates.</p> <p>See “Configuring alternate client backup” on page 74.</p>

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 2003. This method is for Windows 2003 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.

Before running the alternate client backup

Your volume configuration must be prepared and tested for the snapshot method you plan to use.

See the topics on software-based snapshot methods:

See “Software-based snapshot methods” on page 123.

Example alternate client backup configurations

Consider the following example configurations:

Table 3-3 Alternate client backup configurations

Configuration	Description
Client data is on an EMC disk array in split-mirror mode	To run the backup on an alternate client: choose Standard as the policy type, select Perform snapshot backups , Perform off-host backup , and Use alternate client . Then select the alternate client. On the Snapshot Options display, specify an EMC TimeFinder snapshot method. If the data is in an Oracle database, select Oracle as the policy type.
Client data is replicated on a remote host	To run the backup on the replication host (alternate client), choose: Standard as the policy type, select Perform snapshot backups , Perform off-host backup , and Use alternate client . Then select the alternate client (the replication host). On the Snapshot Options display, specify the VVR snapshot method.

Table 3-3 Alternate client backup configurations *(continued)*

Configuration	Description
Client data is on a JBOD array in VxVM volumes with snapshot mirrors configured	<p>To run the backup on the alternate client, choose: Standard (for UNIX client) or MS-Windows (Windows client) as the policy type and Perform snapshot backups, Perform off-host backup, and Use alternate client. Then select the alternate client. On the Snapshot Options display, specify the FlashSnap method.</p> <p>If the client data consists of many files, or to do individual file restore from raw partition backups: select FlashBackup or FlashBackup-Windows as the policy type.</p> <p>Note: Other combinations of policy type and snapshot method are possible, depending on many factors: your hardware configuration, file system and volume manager configuration, and installed NetBackup product licenses.</p>

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.
- On the other hand, a raw partition can be specified in its usual symbolic-link form (such as `/dev/rdisk/c0t1d0s1`). Do not specify the actual device name that `/dev/rdisk/c0t1d0s1` points to. For raw partitions, Snapshot Client automatically resolves the symbolic link to the actual device.
- 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.
- On Windows, the `\` must be entered in the **Backup Selections** list after the drive letter (for example, `D:\`).
See "Configuring a FlashBackup policy" on page 83.

Multiple data streams

Multiplexing is not supported for the Third-Party Copy Device method. (Keep in mind that multiplexing is not the same as 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 incremental backup of mirror-based snapshots

For incremental backup of a mirror type snapshot, note the following issues.

Incremental backup options (Windows only)

NetBackup provides two incremental backup options for Windows clients:

- Based on timestamps
- Based on archive bit (for Windows clients only)


These options are available in the NetBackup Administration Console under **Host Properties** for the client. Right-click the client and select **Properties**. The Client Properties dialog box is displayed. Click **Windows Client > Client Properties**. These options are described in the NetBackup Administrator's Guide, Volume I.

If the snapshot is created on a mirror, such as a Volume Manager volume using VxVM, archive bit incremental backups cannot be used. Archive bit incremental backups are incompatible with snapshots on mirrors. For incremental backups, select **Based on time-stamps**.

Snapshot sequence and access time

After a mirror snapshot backup, the access time of the files in the snapshot is reset to the access time of the original (primary) disk data. The access time of the snapshot data continues to match that of the data on the primary until the mirror is resynchronized with the primary. This behavior is normal for Snapshot Client and applies to primary client backups as well as to off-host backups.

Figure 3-1 Simplified View of Snapshot Sequence And Access Time



The figure shows the following phases:

Phase	Snapshot sequence and access time description
Phase 1	Last file access on primary: 08/29/05 at 8:01 pm
Phase 2	Mirror was synchronized with primary and split at 8:24 pm.
Phase 3	Snapshot taken on mirror at 8:24 pm.
Phase 4	Backup of mirror was completed at 10:14 pm; file access time on mirror is reset to 8:01 pm.

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.

Two different products in NetBackup perform snapshots: Snapshot Client (described in this guide) and Windows Open File Backups (described in the NetBackup Administrator's Guide, Volume I). These products operate independently of each other. Disabling one does not disable the other.

Disabling Open File Backups on Windows

Do the following steps in the NetBackup Administration Console.

To disable Open File Backup snapshots for individual clients

- 1 Click **Host Properties > Master Servers** > double-click master server > **Client Attributes**.
- 2 Click the **Windows Open File Backup** tab.
- 3 Select a client that is enabled for open file backups.
- 4 Clear the check box for **Enable Windows Open File Backups for this client**.
- 5 Click **Apply**.

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 Client** 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**.

Note: Note the following regarding FlashBackup

If the policy type is FlashBackup-Windows, select a non-FlashBackup policy type, such as MS-Windows, which clears the **Perform snapshot backups** check box.

If the policy type is FlashBackup and the policy was configured to use the snapctl driver for Solaris clients or VxFS snap driver for HP clients (with CACHE= entries in the **Backup Selections** tab): select a non-FlashBackup policy type, such as **Standard**. Otherwise, the policy continues to create snapshots even though the **Perform snapshot backups** check box is cleared.

Backups of raw partitions using policy types other than FlashBackup or FlashBackup-Windows do not allow individual file restore. Individual file restore from a raw partition backup is a feature of FlashBackup.

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 82.
- 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 7.x 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.

Note: Earlier versions of WSFC (such as those versions that were shipped with Windows versions before Windows 2003 SP1) do not allow extended maintenance mode functionality. If the cluster does not support placing disks in extended maintenance mode, it is still possible to perform raw restores to an alternate, non-shared disk.

- 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**.

The screenshot shows the 'Change Policy - Flash_Backup_policy1' dialog box. The 'Server' field is set to 'nbu-id1'. The 'Attributes' tab is selected, showing the following configuration:

- 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 (with an 'Options...' button)
 - ☐ 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)

At the bottom right, there are buttons for 'OK', 'Cancel', and '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 60.

- 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 127.
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/clt0d0s6
                     /dev/vx/rdisk/volgrp1/voll
```

HP examples `/dev/rdisk/clt0d0`
 `/dev/vx/rdisk/volgrp1/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/rdisk/clt0d0s6`. **Note** the `r` in `/rdsk`.

Note: Wildcards (such as `/dev/rdisk/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/rdisk/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: `snaptcl` 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 `snaptcl` driver (Solaris) and VxFS snap driver (HP).
- The cache partition for the `snaptcl` 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
```


Configuring FlashBackup policy for backward compatibility (UNIX only)

```
/dev/rdisk/clt4d0s3
/dev/rdisk/clt4d0s4
```


- See “Requirements for the cache partition” on page 90.
- See “About the cache partition” on page 89.

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 4-1 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 `cctxdxsx`.

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 125.

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 90.
- `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 4-1 shows a **Backup Selections** list that generates four backups:

Table 4-1 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 4-1:

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
- About configuring VxVM
- Modifying the VxVM or FlashSnap resync options for point in time rollback
- 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:

- Supports NetBackup clients on Solaris, HP, AIX, Linux, and Windows. The master server can be on any supported operating system.

- 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.
- Enables block-level restore and file promotion from VxFS_Checkpoint snapshots (UNIX) and file promotion from NAS_Snapshot. Also enables Fast File Resync from VxVM and FlashSnap snapshots on Windows.
- Enables rollback from the backups that were created using the following: VxFS_Checkpoint, VxVM, FlashSnap, NAS_Snapshot, or disk array methods.
- Can restore to a different path or host.
- Provides the resource management by means of a rotation schedule.
- Supports Oracle, Microsoft Exchange, DB2, SAP, and SQL-Server databases.

Instant Recovery requirements

Note the following requirements:

- For snapshots using Storage Checkpoints, by means of NetBackup's VxFS_Checkpoint method, note that Solaris clients must have VxFS 3.4 or later (HP clients VxFS 3.5 or later, Linux and AIX clients VxFS 4.0 or later) with the Storage Checkpoint feature.
- For VxVM snapshot volumes on UNIX, clients must have VxVM 3.2 or later with the FastResync feature. Linux and AIX clients require VxVM 4.0 or later. Windows clients must have Storage Foundations for Windows version 3.1 or later.
- For Instant Recovery with DB2, Oracle, Exchange, SAP, or SQL-Server databases, refer to the appropriate NetBackup database agent guide.
- For replication hosts (using NetBackup's VVR method), clients must have VxVM 3.2 or later with the Veritas Volume Replicator feature. Linux and AIX clients require VxVM 4.0 or later.

Instant Recovery restrictions

Note the following restrictions:

- For snapshots using Storage Checkpoints, Instant Recovery supports file systems with the Version 4 disk layout or later. Older disk layouts must be upgraded to Version 4 or later.

- No-data Storage Checkpoints (those containing file system metadata only) are not supported.
- Instant Recovery snapshots must not be manually removed or renamed, otherwise the data cannot be restored.
- Instant Recovery does not support the VxVM, FlashSnap, and VVR snapshot methods when used with VxVM volume sets.
- On Linux, Instant Recovery is not supported by disk-array based snapshot methods.
- For Instant Recovery backups of data that is configured on VxVM volumes on Windows, the VxVM volume names must be 12 characters or fewer. Otherwise, the backup fails.
- 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 96.

- Instant Recovery restores can fail from a backup that a FlashSnap off-host backup policy made.
From a policy that was configured with the **FlashSnap off-host backup** method and with **Retain snapshots for Instant Recovery** enabled, the backups that were made at different times may create snapshot disk groups with the same name. As a result, only one snapshot can be retained at a time. In addition, NetBackup may not be able to remove the catalog images for the snapshots that have expired and been deleted. It appears that you can browse the expired snapshots and restore files from them. But the snapshots no longer exist, and the restore fails with status 5.

- For Instant Recovery, Veritas recommends that a primary volume be backed up by a single Instant Recovery policy. If the same volume is backed up by two or more Instant Recovery policies, conflicts between the policies may occur during snapshot rotation. Data loss could result if the policies are configured for snapshots only (if the policies do not back up the snapshots to separate storage devices).

Consider the following example: Two policies use the same snapshot device (or VxFS storage checkpoint) to keep Instant Recovery snapshots of volume_1.

- Instant Recovery policy_A creates a snapshot of volume_1 on the designated snapshot device or storage checkpoint.
- When Instant Recovery policy_B runs, it removes the snapshot made by policy_A from the snapshot device or storage checkpoint. It then creates its own snapshot of volume_1 on the snapshot device or storage checkpoint. The snapshot created by policy_A is gone.

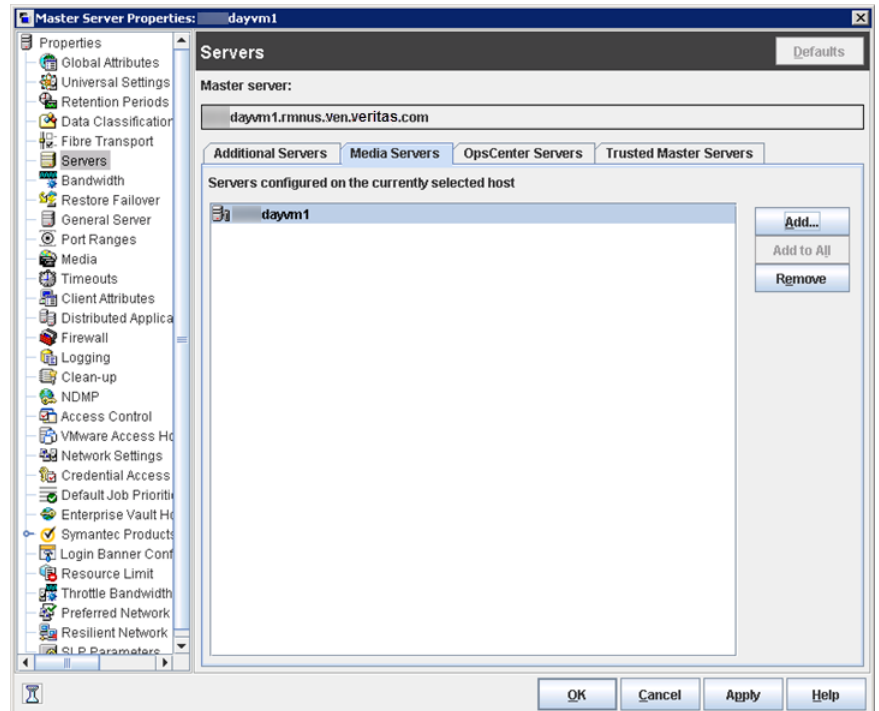
Note: Even if each policy has its own separate snapshot devices, conflicts can occur when you browse for restore. Among the available snapshots, it may be difficult to identify the correct snapshot to be restored. It is therefore best to configure only one policy to protect a given volume when you use the Instant Recovery feature of NetBackup.

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**.



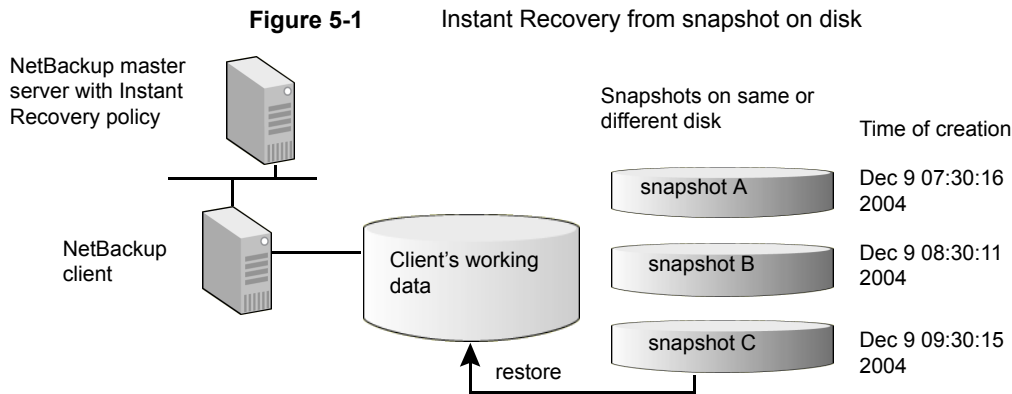
- 2 Make sure that the media server is listed under **Additional Servers**, not under **Media Servers**.

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 5-1 shows an example.



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

An Instant Recovery backup creates a snapshot on disk and optionally backs up the client's data to a storage device. The location of the snapshot depends on the type of snapshot method that is configured in the policy.

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. For many disk array snapshot methods, you can also specify the particular devices on which to retain the snapshots, and the order in which they are used.


About snapshot rotation

NetBackup Snapshot Client implements a rotation scheme for managing snapshot resources for some snapshot methods.

In the following bulleted example, the system is configured to have three snapshots available for instant restore ("undo") at any time:

- The state of the client's backup selection list is captured in the first Instant Recovery snapshot, and the snapshot is retained. In other words, it is not deleted, even if the backup image is also copied to a storage unit.
- When the backup policy runs again, recent changes in the client's backup selection list are captured in a second snapshot. This snapshot is created on a second device that was allocated for this purpose. The first snapshot is retained, in case a restore is needed from that data.
- When the backup policy runs a third time, a third snapshot is created on the third allocated device. At this point three snapshots representing three different states of the client data are available for restore.
- When the policy runs a fourth time, no additional snapshot device is available: one of the existing devices must be reused. The system "rotates" the devices, overwriting the first snapshot with the fourth (most recent) one. Although the earliest snapshot is now gone, the three most recent snapshots are available for restore.

Figure 5-2 Snapshot rotation, for multiple restore (undo) levels



In this figure, the next Instant Recovery backup overwrites the snapshot that was made at 12:00 noon.

Means of controlling snapshots

Depending on the snapshot method you have selected for the policy, one of two means of managing snapshots is available: the **Snapshot Resources** pane or the **Maximum Snapshots** (Instant Recovery only) parameter.

Snapshot Resources pane

The **Snapshot Resources** pane (on the policy's **Snapshot Options** dialog box) is available only for certain disk array snapshot methods on UNIX clients.

The **Snapshot Resources** pane provides the most control over Instant Recovery snapshots. You can specify which devices to use for the snapshots, and the order in which the devices are used. Click the **Add** option to display the following:

Figure 5-3 Example of Add Snapshot Resource dialog

The dialog box is titled 'Add Snapshot Resource'. It contains three input fields: 'Array Serial #' with the value '0001 36268344', 'Source Device :' with the value '01 41', and 'Snapshot Device(s) :' with the value '0122;0123;0124'. At the bottom are three buttons: 'OK', 'Cancel', and 'Help'.

- The serial number of the array is specified in the **Array Serial #** field. Contact your array administrator to obtain the disk array serial numbers and designators (unique IDs) for the array.

The unique ID snapshot resource or source LUN containing the primary data is specified in the **Source Device**.

The maximum number of snapshots to retain is determined by the number of configured devices in the **Snapshot Device(s)** field. For example, if you enter two devices, only two snapshots can be retained. The above example specifies three devices (0122;0123;0124), so three snapshots can be retained. When the maximum is reached, the fourth snapshot overwrites the first one.

- The particular devices to use for the snapshots are those named in the **Snapshot Device(s)** field.
- The order in which the devices are listed in the **Snapshot Device(s)** field determines their order of use. Device 0122 is used for the first snapshot, 0123 for the second, and 0124 for the third.

Preconfiguration of the snapshot devices may be required.

See the appropriate topic for your disk array and snapshot method.

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. This parameter is available only for certain snapshot methods. Unlike the **Snapshot Resources** pane, it cannot specify which devices to use or the order of their use. (The **Maximum Snapshots** parameter and the **Snapshot Resources** pane are mutually exclusive.)

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.

Note: For Windows clients using the VSS method on disk arrays that are configured for clones or mirrors: you must synchronize the clones or mirrors with their source before you run the backup.

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
- For the policy type, select **Standard**, **MS-Windows**, **FlashBackup-Windows**, or the database agent type appropriate for the client(s).
- 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
- Click the **Options** option to select the snapshot method.

For a new policy, you can skip this step to let NetBackup select the method (auto is the default method).

-
- Select a snapshot method from the pull-down list. For creating an Instant Recovery snapshot, the available methods are:

auto

(UNIX or Windows)

NetBackup selects the snapshot method. If auto is the method, only one snapshot can be retained at a time.

FlashSnap

(UNIX or Windows)

Uses the VxVM FlashSnap feature and VxVM FastResync to create the snapshot. VxVM mirrors must be configured:

See “About configuring VxVM” on page 106.

This method can also use instant snapshots:

See “About VxVM instant snapshots” on page 132.

For Windows clients, FlashSnap can use the Fast File Resync feature of Storage Foundation for Windows 4.1.

NAS_Snapshot (UNIX or Windows)	<p>Uses the NDMP V4 snapshot extension to create the snapshot on the NAS-attached disk.</p> <p>See “Setting up a policy for NAS snapshots” on page 120.</p>
OST_FIM	<p>The name of the snapshot method that is selected in a policy configured for snapshot replication using Replication Director. The name represents OpenStorage Frozen Image Method.</p> <p>Refer to the NetBackup Replication Director Solutions Guide for more details.</p>
VSS (Windows 2003 only)	<p>VSS selects the actual snapshot method depending on which snapshot provider is configured on the client. For example, if the client data is stored on an EMC CLARiiON disk array, and the array administrator has fully configured the array and its snapshot capability: VSS selects the appropriate EMC CLARiiON snapshot method for the backup.</p>
VxFS_Checkpoint (UNIX only)	<p>Uses VxFS Storage Checkpoint to create the snapshot. A new Storage Checkpoint is created whenever a backup using the VxFS_Checkpoint method is executed.</p>
VxVM (UNIX or Windows)	<p>Uses VxVM FastResync to create the snapshot. VxVM mirrors must be configured:</p> <p>See “About configuring VxVM” on page 106.</p> <p>The VxVM method can also use instant snapshots: See “About VxVM instant snapshots” on page 132.</p> <p>For Windows clients, VxVM can use the Fast File Resync feature of Storage Foundation for Windows 4.1.</p>
VVR (UNIX)	<p>Creates a snapshot of a VxVM volume on a Veritas Volume Replication host.</p>
The disk array snapshot methods.	<p>See the appropriate topic for your disk array and snapshot method.</p>

- Change parameter values for the method, if needed.
- When you are finished, click **OK**.

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. This option is required for the **NAS_Snapshot** method. 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 the snapshot uses **VxFS_Checkpoint** or is a **VxVM** space-optimized snapshot, it is created on the same device as the device containing the original data. In that case, you may want to create another policy to back up the data to a separate device.
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.
 - Snapshot Client policies do not support the **ALL_LOCAL_DRIVES** entry in the policy's **Backup Selections** list, except for the policies that are configured with the VMware method.
- 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.

Use the following copy-on-write snapshot methods with Instant Recovery only when small amounts of data (individual files) are likely to be restored:

- VxVM space optimized instant snapshots
- VxFS_Checkpoint
- The following copy-on-write disk array methods if the snapshot is not fully allocated:
 - EMC_CLARiiON_SnapView_Snapshot
 - EMC_TimeFinder_Snap
 - HP_EVA_Snapshot
 - HP_EVA_Vsnap
 - Hitachi_CopyOnWrite

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.

Large restores from an Instant Recovery snapshot

For restore of an entire raw partition or large file system, Veritas recommends a fully allocated snapshot method.

Examples are the following:

- VxVM as a snapshot mirror (not space-optimized)
- EMC_CLARiiON_Snapview_Clone
- EMC_TimeFinder_Clone
- EMC_TimeFinder_Mirror
- HP_EVA_Snapclone
- Hitachi_ShadowImage
- IBM_DiskStorage_FlashCopy
- IBM_StorageManager_FlashCopy
- Any copy-on-write disk array method that is configured with a fully-allocated snapshot.

About configuring VxVM

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.

Before you run an Instant Recovery policy to back up VxVM volumes, one or more mirrors must be created. The primary volumes must be enabled for FastResync. Note that on Windows, FastResync is enabled by default.

Creating a snapshot mirror

You can create a snapshot mirror using VxVM commands.

To create a snapshot mirror on Windows

- 1 For a volume that is associated with a drive letter, enter:

```
vxassist snapstart X:
```

where X is the drive letter. This command creates a snapshot mirror of the designated drive.

- 2 For a volume that is not associated with a drive letter, enter:

```
vxldg -g disk_group dginfo
```

This command shows information for the specified disk group, including the names of the volumes that are configured for that group.

- Create the snapshot by entering the following:

```
vxassist snapstart \Device\HarddiskDmVolumes\disk_group\Volume_name
```

This command creates a snapshot mirror of the designated Windows volume.

To create a snapshot mirror on UNIX

- 1 Add dco (data change object) logs to the primary volume:

```
/usr/sbin/vxassist -g disk_group addlog volume_name logtype=dco
```

- 2 Enable FastResync on the volume:

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

- 3 Prepare a new mirror for the Instant Recovery snapshot:

- Create the mirror:

```
/usr/sbin/vxassist -g disk_group snapstart primary_volume
```

Wait until the mirror is synchronized (status SNAPDONE, or **State** field reads **Ready** in the volume's properties display).

- To verify that the mirror is synchronized, enter:

```
/usr/sbin/vxprint -g disk_group -q -t -e 'assoc="primary_volume"'
```

About creating Instant Snapshots

The following procedures are required for Instant Recovery backups when using the full-sized or space-optimized instant snapshot options in VxVM 4.0. For Instant Recovery backups, VxVM 4.0 or later instant snapshots are supported by Snapshot Client's FlashSnap, VxVM, and VVR methods.

Creating space-optimized snapshots

A cache object called NBU_CACHE must be created in the disk group containing the volume to be backed up. NetBackup recognizes the cache object and uses it to create a space-optimized snapshot.

To create space-optimized snapshots

1 Create the parent volume:

```
/usr/sbin/vxassist -g disk_group make volume size layout=layout
logtype=dco dconversion=20 [drl=no|sequential|yes]
[ndcomirror=number] fastresync=on
```

Where:

- Brackets [] indicate optional items.
- *make volume* specifies the name of the volume snapshot.

2 Create the cache object:

```
/usr/sbin/vxassist -g disk_group make cache_volume size
layout=layout init=active
```

3 Label the cache object:

```
/usr/sbin/vxmake -g disk_group cache NBU_CACHE
cachevolname=cache_volume
```

4 Enable the cache object:

```
/usr/sbin/vxcache -g disk_group start NBU_CACHE
```

5 Take the initial snapshot:

```
/usr/sbin/vxsnap -g disk_group make
source=volume/newvol

=SNAP_vol1_NBU/cache=NBU_CACHE
```

Creating full-sized snapshots

Unlike the space-optimized snapshot, NetBackup cannot create VxVM full-sized instant snapshots: you must create them before running the backup, as explained in the following procedure. You must create one full-sized instant snapshot for each backup you want to run.

To create full-sized snapshots

- 1 Enter the following to create the parent volume:

```
/usr/sbin/vxassist -g disk_group make volume length layout=layout
logtype=dco dconversion=20 [drl=no|sequential|yes]
[ndcomirror=number] fastresync=on
```

Where:

- Brackets [] indicate optional items.
- `make volume` specifies the name of the volume snapshot.

- 2 Create a volume for a full-sized instant snapshot:

- Determine the required size for the snapshot volume:

```
# LEN='vxprint -g disk_group -F%len volume'
```

- Find the name of the DCO volume:

```
# DCOVOL='vxprint -g disk_group -F%dconame volume'
```

- Discover the DCO volume's region size (in blocks):

```
# RSZ='vxprint -g disk_group -F%regionsz $DCOVOL'
```

- Create a volume that is named `voluname_NBU`, of the required size and redundancy.

The volume name must end with `_NBU`. In the following example, the volume is named `SNAP_voll_NBU`.

```
vxassist -g disk_group make SNAP_voll_NBU $LEN layout=mirror
nmirror=number logtype=dco drl=no dconversion=20
ndcomirror=number regionsz=$RSZ init=none
[storage attributes ...]
```

The number for `nmirror` should equal the number for `ndcomirror`.

Note: For Linux, the init value should be `init=active` instead of `init=none`.

For Solaris 10 with Storage Foundation 5.1, the init value should be `init=active` instead of `init=none`.

- Create the mirror:

```
vxsnap -g disk_group make source=volume/snapvol=SNAP_voll_NBU/syncing=on
```

- 3 Set the **Maximum Snapshots (Instant Recovery only)** value on the NetBackup **Snapshot Client Options** dialog.

Using the VxVM 3.5 GUI to configure VxVM mirrors

The following are the steps for using the VxVM 3.5 graphical user interface to configure VxVM mirrors for Instant Recovery backups.

To use the VxVM 3.5 GUI to configure VxVM mirrors

- 1 Make sure that FastResync is enabled on the primary VxVM volume.
 - In the VEA console, right-click the volume and click **Properties** from the pop-up menu.
 - The **FastResync** field states whether or not FastResync is enabled.
 - Click **Cancel**.
 - If FastResync is disabled, right-click the volume again and select **Fast Resync > Add** from the pop-up menu to enable it.
- 2 Prepare a new mirror for Instant Recovery backup.
 - Create a new mirror by right-clicking the volume and selecting **Snap > Snap Start**.
 - Make sure FastResync is enabled on the mirror. Click **OK** to create the mirror and start full synchronization.
- 3 On the **Mirrors** tab, ensure that synchronization has completed as indicated by **Snap Ready** in the **Status** field.

Modifying the VxVM or FlashSnap resync options for point in time rollback

Several options can be set for Instant Recovery point in time rollback when multiple volumes are involved in the rollback. For better performance, you can modify these options to suit the circumstances of the rollback.

The following options can be modified for rollback of a VxVM or FlashSnap snapshot:

- Maximum number of volumes to resynchronize concurrently
- Sync region size in MB (UNIX only)
- Sync I/O delay in milliseconds (UNIX only)

By default, these resync options have the same values for restore as they had for the backup. For the defaults, see the following section:

See “Configuration parameters for Snapshot Client” on page 65.

To modify the VxVM or FlashSnap resync options for point in time rollback

- 1 Create the following file:

```
/usr/opensv/netbackup/SYNC_PARAMS
```

- 2 In the file, enter the numeric values for the options, on one line. The numbers apply to the options in the bulleted list above, in that order.

For example:

```
6 3 1000
```

This example resets the options as follows:

- Maximum number of volumes to resynchronize concurrently = 6
- Sync region size in MB (UNIX only) = 3
- Sync I/O delay in milliseconds (UNIX only) = 1000

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.

The same is true for the **Snapshot Resources** pane on the **Snapshot Options** dialog box. If the snapshot method requires snapshot resources, the maximum number of snapshots is determined by the number of devices that are specified in the **Snapshot Device(s)** field. For example, if two devices are specified, only two snapshots can be retained at a time. Either the **Snapshot Device(s)** field or the snapshot retention period in the lifecycle policy can determine the retention period.

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 8.0 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 polices 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: tashinall_1204305543 cannot be used with a
lifecycle policy NoSnapshot that does not include a snapshot
destination.
```

Error 156 can be a result of different problems, listed below are some of them:

VxVM failing to get the version of the disk group, run appropriate VxVM command outside of NetBackup to see whether you can get the version information for the disk group in use.

Bpfis log

```
10:43:58.436 [28336] <32> onlfi_fim_dev_info: FTL - VfMS error 11;
see following messages:10:43:58.437 [28336] <32> onlfi_fim_dev_info:
FTL - Fatal method error was reported
10:43:58.437 [28336] <32> onlfi_fim_dev_info:
FTL - vfm_freeze: method: vxvm, type: FIM, function: vxvm_freeze
10:43:58.437 [28336] <32> onlfi_fim_dev_info: FTL - VfMS method error 10;
see following message:
10:43:58.437 [28336] <32> onlfi_fim_dev_info:
FTL - vxvm__get_dgversion: Cannot get version for disk group: dgdb001
10:43:58.437 [28336] <4> onlfi_thaw:
INF - Thawing /ora/db001/data001 using snapshot method vxvm.
10:43:58.448 [28336] <4> onlfi_thaw: INF - do_thaw return value: 0
10:43:58.454 [28336] <16> bpfis:
FTL - snapshot preparation failed, status 156
```

The device that is to backup by this process is being used by another process. Check whether any other process is holding the same device.

Bpfis log

```
00:26:19.025 [2826] <2> onlfi_vfms_logf: INF - lock pid(2902) != pid(2826):
/usr/openv/netbackup/online_util/db_cntl/___LOCKFILE_EMC:
SYMMETRIX:970960001000
00:26:19.025 [2826] <2> onlfi_vfms_logf: INF - TimeFinder_rebuild:
Cannot get lock on device: /dev/rdisk/c3t5006048C4A85A400d1s2 .....
00:26:19.025 [2826] <32> rebuild_fim_list: FTL - TimeFinder_rebuild:
Cannot get lock on device: /dev/rdisk/c3t5006048C4A85A400d1s2
00:26:19.025 [2826] <32> splthost_rebuild: FTL - rebuild_fim_list() failed
00:26:19.037 [2826] <4> bpfis Exit:
INF - EXIT STATUS 156: snapshot error encountered
```

Policy validation fails for valid backup selection. If there filer's volume is mounted on a windows client, run NetBackup client service on the client and the alternate client with a valid credentials to access CIFS share, and check that the filers are up, and the volume is seen mounted on the windows client.

Bpfis log

```
11:49:40.727 [15240.13716] <16> bpfis main: FTL - process_fs_list() failed,
status 71
11:49:40.727 [15240.13716] <2> ol_cleanup:
INF - removing
C:\Program Files\Veritas\NetBackup\temp\unknown+15240+1.std_filelist
11:49:40.727 [15240.13716] <4> bpfis Exit:
INF - EXIT STATUS 71: none of the files in the file list exist
11:49:40.743 [15240.13716] <2> stop_keep_alive_thread:
INF - Stop keep_alive thread
11:49:40.743 [15240.13716] <2> bpfis Exit: INF - Close of stdout
```

For windows client, live browse from the snapshot fails with the following error message. Make sure that NetBackup client service on the client and the alternate client is running with a valid credential to access CIFS share

ERROR: permissions denied by client during rcmd.

Snapshot backup for windows client fail with status 55. Make sure that NetBackup client service on the client and the alternate client is running with a valid credential to access CIFS share.

Bpfis log

```
10:46:00.131 [2612.7880] <4> bpcd_request_mount:
get volume guid from <frag-id> failed with 55
```

```
10:46:00.131 [2612.7880] <4> bpcd_request_mount:
return mntdev NO_MOUNT_DEVICE
10:46:00.131 [2612.7880] <16> bpcd_get_fileinfo:
bpcd_request_mount failed with error[55]
10:46:00.147 [2612.7880] <2> bpcd_get_fileinfo: sent status 55 to bpdbm
```

Live browse or 'backup from snapshot' operation for windows client fail with error 43, status 156. Enable create_ucode & convert_ucode on primary volume.

Bpfis log

```
04:43:44.656 [3040.3900] <2> onlfi_vfms_logf:
INF - snapshot services:ostfi:Wed Aug 24 2011
04:43:44.640000 <Thread id - 3900> Failed to import snapshot [*****]
04:43:44.718 [3040.3900] <2> onlfi_vfms_logf:
INF - snapshot services: ostfi:Wed Aug 24 2011
04:43:44.718000 <Thread id - 3900> OST Library call failed with message
(STS API sts_create_export failed with
error code : 2060022)
04:43:44.718 [3040.3900] <2> onlfi_vfms_logf: INF - snapshot services:
ostfi:Wed Aug 24 2011
04:43:44.718000 <Thread id - 3900> COSTPlugin::importTreeNode -
Could not import device[*****]
```

NBUAdapter log

```
0 RESTORE :2104 111 0 115851 2011/08/24 04:58:51
Volume name = f3070-238-15:/NetBackup_1314174973_mirror,
dest snap f270-247-156_test6
1 RESTORE :2104 44 0 115851 2011/08/24 04:58:51
share name:NBU_Share_NetBackup_1314174973_mirror_
f270-247-156_test6_2011_08_24_04_36_14
0 RESTORE :2104 6 0 115851 2011/08/24 04:58:51 add_cifs_export :
Failed to get response. Error 22 : Directory
"/vol/NetBackup_1314174973_mirror/.snapshot/f270-247-156_test6"
does not exist.
1 STRWIDE :2104 1 0 115851 LFB:
Need 111 chars to store wide copy of UTF8
'Directory "/vol/NetBackup_1314174973_mirror/.snapshot/f270-247-156_test6"
does not exist. `
```

Network Attached Storage (NAS) snapshot configuration

This chapter includes the following topics:

- About NAS snapshot overview
- Notes on NAS_Snapshot
- Logging on to the NetBackup Client Service as the Administrator
- Setting up a policy for NAS snapshots
- NAS snapshot naming scheme

About NAS snapshot overview

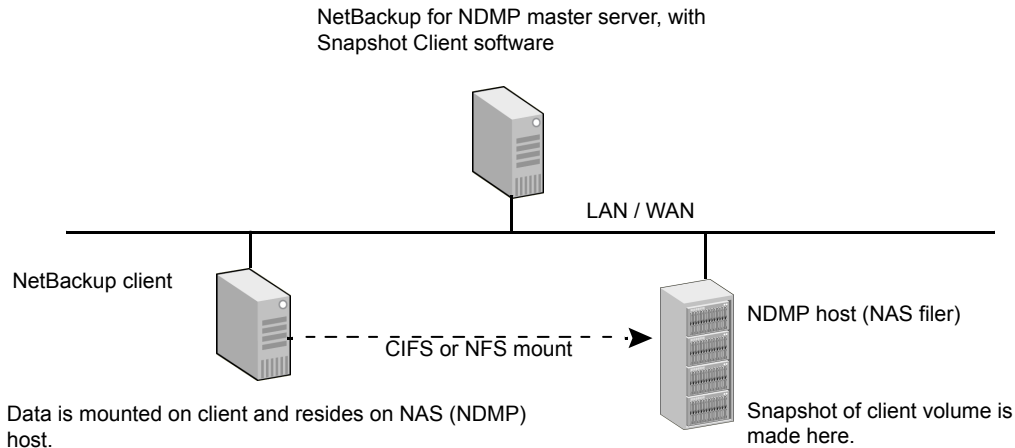
By means of Snapshot Client and the NDMP V4 snapshot extension, NetBackup can make snapshots of client data on a NAS (NDMP) host. The client data must reside on the NAS host and be accessible on the client through an NFS mount on UNIX or CIFS on Windows.

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.

Note: NetBackup for NDMP software is required in addition to Snapshot Client.

See the following diagram for an overview.

Figure 6-1 NAS snapshot environment



In the NetBackup policy, enter the following:

- For Windows client:
`\\ndmp_hostname\share_name`
- 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.

Notes on NAS_Snapshot

The following notes apply to backups that are made with the NAS_Snapshot method:

- Snapshots of NAS host data are supported for NetBackup clients running Windows (32-bit system and 64-bit system), Solaris, Linux, and AIX.
- Note these software requirements and licensing requirements:
 - On the NetBackup server, both NetBackup for NDMP and Snapshot Client software must be installed and licensed. In addition, a NetBackup for NDMP license must be purchased for each NDMP host (filer).
 - NetBackup clients that are used to perform backups must have Snapshot Client installed.

- On NetBackup clients for Oracle: NetBackup for Oracle database agent software must be installed on all clients.
- The NAS host must support NDMP protocol version V4 and the NDMP V4 snapshot extension, with additional changes made to the snapshot extension. The *NetBackup Snapshot Client Configuration* online pdf contains a list of NAS vendors that NetBackup supports for NAS snapshots. This online pdf includes requirements specific to your NAS vendor.
 See: <http://www.veritas.com/docs/000081320>
- NetBackup must have access to each NAS host on which a NAS snapshot is to be created. To set up this authorization, you can use either of the following:
 - In the NetBackup Administration Console: the **Media and Device Management > Credentials > NDMP Hosts** option or the NetBackup Device Configuration Wizard.
 - OR
 - The following command:


```
tpconfig -add -nh ndmp_host -user_id user_ID -password password
```
- The client data must reside on a NAS host and be mounted on the client by means of NFS on UNIX or CIFS on Windows. For NFS mounts, the data must not be auto-mounted, but must be hard (or manually) mounted.
- For NAS snapshot, you must create a NAS_Snapshot policy.
 See “Setting up a policy for NAS snapshots” on page 120.
- On Windows clients, to restore files from a NAS_Snapshot backup, the NetBackup Client Service must be logged in as the Administrator account. The NetBackup Client Service must not be logged in as the local system account. The Administrator account allows NetBackup to view the directories on the NDMP host to which the data is to be restored. If you attempt to restore files from a NAS_Snapshot and the NetBackup Client Service is logged in as the local system account, the restore fails.
 See “Logging on to the NetBackup Client Service as the Administrator” on page 119.

Logging on to the NetBackup Client Service as the Administrator

Use the following procedure.

To log on to the NetBackup Client Service as the Administrator

- 1 In Windows Services, double-click the NetBackup Client Service.
- 2 Check the **Log On** tab: if the service is not logged on as Administrator, stop the service.
- 3 Change the log in to the Administrator account and restart the service.
- 4 Retry the restore.

Setting up a policy for NAS snapshots

This section explains how to set up a policy for making snapshots of NAS data.

To set up a policy for NAS snapshots

- 1 Open the NetBackup Administration Console on the NetBackup for NDMP server.
- 2 Click **NetBackup Management > Policies** in the left pane.
- 3 In the **All Policies** pane, double-click an existing policy or right-click and create a new one.
- 4 For **Policy type**: select Standard for UNIX clients, MS-Windows for Windows clients, SAP for UNIX clients, or Oracle for UNIX clients that are configured in an Oracle database.
- 5 For Storage unit, select **Any_available** if this policy is for a NAS snapshot. Note the following:
 - Although the policy cannot execute without a specified storage unit, NetBackup does not use the storage unit. The snapshot is created on disk regardless of which storage unit you select. Note that the storage unit you select is not reserved, so it is free for use by other policies.
 - For Oracle policies, the policy uses the storage unit you specify, but only for backing up archive logs and control files.
- 6 Select **Perform snapshot backups** and **Retain snapshot for Instant Recovery or SLP management**.
- 7 Select the **Perform off-host backup** and **Use** options.

- 8 Select **Data Mover** from the **Use** list and **Network Attached Storage** from the **Machine** list.

When the policy executes, NetBackup automatically selects the NAS_Snapshot method for creating the snapshot.

As an alternative, you can manually select the NAS_Snapshot method from the **Options** dialog from the policy **Attributes** display.

- 9 On the **Schedule Attributes** tab, select the following:

- **Instant Recovery**

Choose **Snapshots only**. The other option (**Snapshots and copy snapshots to a storage unit**) does not apply to NAS_Snapshot.

- **Override policy storage unit**

If the correct storage unit was not selected on the **Attributes** tab, select it here.

- 10 For the **Backup Selections** list, specify the directories, volumes, or files from the client perspective, not from the NDMP host perspective. For example:

- On a UNIX client, if the data resides in `/vol/vol1` on the NDMP host `nas1`, and is NFS mounted to `/mnt2/home` on the UNIX client: specify `/mnt2/home` in the **Backup Selections** list.
- On a Windows client, if the data resides in `/vol/vol1` on the NDMP host `nas1`, and is shared by means of CIFS as `vol1` on the Windows client, specify `\\nas1\vol1`.
- Windows path names must use the Universal Naming Convention (UNC), in the form `\\server_name\share_name`.
- The client data must reside on a NAS host. The data must be mounted on the client by means of NFS on UNIX or shared by means of CIFS on Windows. For NFS mounts, the data must be manually mounted by means of the `mount` command, not auto-mounted.
- For a client in the policy, all paths must be valid, or the backup fails.
- The `ALL_LOCAL_DRIVES` entry is not allowed in the **Backup Selections** list.

- 11 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.

NAS snapshot naming scheme

The format of a NAS snapshot name is as follows:

NAS+NBU+PFI+client_name+policy_name+sr+volume_name+date_time_string

Note the following:

- The snapshot name always begins with `NAS+NBU+PFI+`
- The plus sign (+) separates name components.
- NAS snapshots reside on the NDMP host (NAS filer).

For example:

`NAS+NBU+PFI+sponge+NAS_snapshot_poll+sr+Vol_15G+2005.05.31.13h41m41s`

Where:

Client name = `sponge`

Policy name = `NAS_snapshot_poll`

`sr` = indicates that the snapshot was created for a NAS snapshot.

Volume name = `Vol_15G`

Date/Time = `2005.05.31.13h41m41s`

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:

See “About array-specific methods vs array-independent methods” on page 146.

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 `snaptcl` driver when you configure FlashBackup in the earlier manner.

See “Configuring FlashBackup policy for backward compatibility (UNIX only)” on page 87.

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/rdisk/diskgroup_1/volume_3`

Or

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

Note: Do not specify wildcards (such as `/dev/rdisk/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 125.

- 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 227.

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 124.

- 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 `snapcachelist` 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 282.

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/rdisk/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/openv/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 `fsckptadm list` command to show Storage Checkpoint disk usage. Refer to the *Veritas File System Administrator's Guide* for more information on `fsckptadm`.

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.

See "About Instant Recovery: block-level restore" on page 234.

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 87.
- 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 124.

See "Entering the cache" on page 127.

- 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 131.
Or refer to your *Veritas Volume Manager* documentation.
- Help is available for configuring a cache object.
See “About VxVM instant snapshots” on page 132.
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 227.

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.

See "Creating space-optimized snapshots" on page 108.

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 28.

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=dco
```

- 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:

```
vxdbg 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 264.

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:

```
vxdg -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:

```
vxdg -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
\snap_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 31.

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 132.

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 137.

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 -R10 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 100.

About VSP

Note: VSP is currently deprecated.

VSP is the Veritas Volume Snapshot Provider, for snapshots of open and active files on NetBackup pre-7.0 Windows clients.

Note: For clients of 7.0 or later versions, NetBackup automatically uses VSS instead of VSP.

For pre-7.0 Windows clients, you can use VSP without Snapshot Client, as explained in the NetBackup Administrator's Guide, Volume I. If the **Busy File Timeout** has expired, no snapshot is created and the backup job may continue without backing

up the busy file. If you use VSP with Snapshot Client, the backup either successfully creates a snapshot of all files, or the backup job fails.

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.

Note: For NetBackup clients (7.0 or later versions) included in a policy that is configured to use VSP, NetBackup automatically uses VSS instead of VSP. In such cases, a message in the detailed status log in the Activity Monitor indicates that VSS is used as the snapshot method.

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 133.

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. (For VxVM snapshot, this support was added in NetBackup 6.5.2.)
- 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 VxVM, this support was added in NetBackup 6.5.2.)
- 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 143.

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:

- About the new disk array snapshot methods
- Disk array configuration tasks
- OS-specific configuration tasks
- About VSS configuration (Windows)
- About EMC CLARiiON arrays
- About EMC Symmetrix arrays
- About HP EVA arrays
- 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

About the new disk array snapshot methods

These topics describe NetBackup's disk array snapshot methods. These methods take advantage of high-speed mirroring and other snapshot capabilities that are provided by particular disk arrays.

The following snapshot methods support only English locale. They do not support I18N (internationalization).

- EMC_CLARiiON_Snapview_Clone
- EMC_CLARiiON_Snapview_Snapshot
- EMC_TimeFinder_Clone
- EMC_TimeFinder_Mirror
- EMC_TimeFinder_Snap
- Hitachi_ShadowImage
- Hitachi_CopyOnWrite
- HP_EVA_Vsnap
- HP_EVA_Snapshot
- HP_EVA_Snapclone
- HP_XP_BusinessCopy
- HP_XP_Snapshot
- IBM_DiskStorage_FlashCopy
- IBM_StorageManager_FlashCopy

About array-specific methods vs array-independent methods

Some snapshot methods are disk array independent. These methods are described in the following topic:

See “Software-based snapshot methods” on page 123.

For these methods, the client platform or the presence of a particular file system or volume manager might determine or limit selection of the snapshot method. However, if the client data resides on an array, the make of the array does not determine the choice of the snapshot method.

The array methods in the current chapter, however, are each written for a particular model of disk array. For instance, the EMC_CLARiiON_SnapView_Clone method is designed for the EMC CLARiiON CX 300/500 and CX3 series arrays. Regardless of client platform or storage stack, you cannot use EMC_CLARiiON_SnapView_Clone on a non-EMC array, or even on a different model of EMC array.

For an up-to-date list of supported disk arrays, refer to the *NetBackup 7.x Snapshot Client Compatibility* document. This document may be accessed from the following link:

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

Advantages of the new array-specific methods

The snapshot methods that are described in this chapter offer advantages over software-based methods. Note the following when choosing between the array-specific methods in this chapter and the snapshot methods that are described in other chapters of this guide.

The new array-specific methods enable a variety of snapshot capabilities that the disk arrays provide:

- All data movement is within the array, saving network bandwidth.
- The Instant Recovery feature of NetBackup Snapshot Client.
See “About Instant Recovery” on page 20.
The legacy array-specific methods (TimeFinder, ShadowImage, BusinessCopy) do not support Instant Recovery.
- Clone and copy-on-write snapshot methods, in addition to full mirror methods (depending on the array). The legacy array methods support only mirrors.
- Automatic provisioning (LUN masking) of snapshot devices on the NetBackup clients. With the legacy array-specific methods, a mirror device had to be manually provisioned before it can be used in a backup. That preconfiguration step is no longer necessary.

About types of disk array methods

Two types of snapshots are designed for each supported disk array. One type creates a full-sized, fully allocated copy (a clone or mirror). Another type uses a copy-on-write system, which is not a full copy and saves disk space. The methods go by various names, depending on the disk array vendor.

All methods are listed in the following topic.

See “Disk array methods at a glance” on page 149.

Note: Some disk array vendors use the term snapshot to refer to a certain kind of point-in-time copy made by the array. In other chapters of this guide, however, snapshot refers more generally to all kinds of point-in-time copies, disk-array based or otherwise. Refer to your array documentation for the definition of array vendor terminology.

Important disk array method notes and restrictions

Note the following:

- The disk array methods support the Veritas File System (VxFS). With a few exceptions, these methods do not support software-based volume managers such as the Veritas Volume Manager (VxVM) or any native Linux or Solaris volume managers. If your client data is configured in Veritas VxVM volumes, use either the legacy array-specific methods for UNIX clients (TimeFinder, ShadowImage, BusinessCopy) or a software-based snapshot method such as VxVM or FlashSnap.

Note: The following array methods support Veritas Volume Manager (VxVM) volumes: Hitachi_CopyOnWrite and Hitachi_ShadowImage. The IBM_DiskStorage_FlashCopy method (on the IBM DS6000) supports VxVM on the AIX platform.

- Use caution with Instant Recovery rollback. An Instant Recovery point-in-time rollback overwrites the entire LUN (source disk) with the contents of the snapshot or mirror disk. If you have multiple file systems or multiple partitions configured on the hardware array LUN (the source disk), one or more of those file systems or partitions sharing the snapshot disk or mirror disk may have older data that you do not want to write back to the source. When the rollback takes place, any older data on the snapshot disk or mirror disk replaces the newer data on the source.
- In a clustered environment, Instant Recovery point-in-time rollback is not supported.
- Except for required preconfiguration as explained in this chapter, do not perform manual operations on any snapshot resources that NetBackup uses. After the preconfiguration is completed as explained in this chapter, NetBackup automatically manages the required LUNs, snapshots, clones, and mirrors.

Warning: If you make other changes to the snapshot resources, the NetBackup catalog may be invalidated. For instance, restores may fail from backups consisting of the snapshots that have been deleted outside the view of NetBackup.

For example, DO NOT do the following:

- Do not remove a snapshot resource that NetBackup created.
- Do not create a snapshot resource in a storage group.

- Do not change the state of a snapshot resource, such as by manually resynchronizing it.
- You should not use array device targets in more than one policy. If you want to reuse devices after destroying an old policy, you need to manually expire the backup images, which exist for that policy.

Disk array methods at a glance

Table 9-1 is an alphabetical listing of the disk array snapshot methods.

For up-to-date information on the arrays and on supported software versions, refer to the *NetBackup Snapshot Client Configuration* document. This document may be accessed from the following link:

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

Table 9-1 Snapshot methods at a glance

Snapshot method	Description and notes
EMC_CLARiiON_Snapview_Clone	For full-volume mirror snapshots with EMC CLARiiON disk arrays with Navisphere. (An EMC CLARiiON clone is actually a full-volume copy mirror, like a Symmetrix BCV.) See “About EMC CLARiiON arrays” on page 162.
EMC_CLARiiON_SnapView_Snapshot	For space-optimized, copy-on-write snapshots with EMC CLARiiON disk arrays with Navisphere. See “About EMC CLARiiON arrays” on page 162.
EMC_TimeFinder_Clone	For full-volume copy (clone) snapshots with EMC disk arrays with Solutions Enabler. See “About EMC Symmetrix arrays” on page 176.
EMC_TimeFinder_Mirror	For full-volume copy (mirror) snapshots with EMC disk arrays with Solutions Enabler. See “About EMC Symmetrix arrays” on page 176.
EMC_TimeFinder_Snap	For space-optimized, copy-on-write snapshots with EMC disk arrays with Solutions Enabler. See “About EMC Symmetrix arrays” on page 176.

Table 9-1 Snapshot methods at a glance (*continued*)

Snapshot method	Description and notes
Hitachi_CopyOnWrite	<p>For space-optimized, copy-on-write snapshots with Hitachi SMS/WMS/AMS, USP/NSC, and USP-V/VM series of arrays.</p> <p>See “About Hitachi SMS/WMS/AMS, USP/NSC, USP-V/VM” on page 201.</p>
Hitachi_ShadowImage	<p>For full-volume copy (mirror) snapshots with Hitachi SMS/WMS/AMS, USP/NSC, and USP-V/VM series of arrays.</p> <p>See “About Hitachi SMS/WMS/AMS, USP/NSC, USP-V/VM” on page 201.</p>
HP_EVA_Snapclone	<p>For full-volume copy (clone) snapshots with Hewlett-Packard EVA disk arrays with CommandView SSSU.</p> <p>Note: this method is the only EVA disk array method that supports Instant Recovery.</p> <p>See “About HP EVA arrays” on page 184.</p>
HP_EVA_Snapshot	<p>For space-optimized, fully allocated copy-on-write snapshots with Hewlett-Packard EVA disk arrays with CommandView SSSU.</p> <p>See “About HP EVA arrays” on page 184.</p>
HP_EVA_Vsnap	<p>For space-optimized, on-demand copy-on-write snapshots with Hewlett-Packard EVA disk arrays with CommandView SSSU.</p> <p>See “About HP EVA arrays” on page 184.</p>
IBM_DiskStorage_FlashCopy	<p>For full-volume copy (clone) snapshots on IBM DS6000 and DS8000 series of arrays with DSCLI version.</p> <p>See “About IBM DS6000 and DS8000 arrays” on page 190.</p>
IBM_StorageManager_FlashCopy	<p>For full-volume copy (clone) snapshots on the IBM DS4000 series of arrays (excluding 4100), with SMcli.</p> <p>See “About IBM DS4000 array” on page 196.</p>

Table 9-1 Snapshot methods at a glance (*continued*)

Snapshot method	Description and notes
HP_XP_BusinessCopy	For mirror-based snapshots with HP-XP arrays See “About HP-XP arrays” on page 206.
HP_XP_Snapshot	For COW-based snapshots with HP-XP arrays See “About HP-XP arrays” on page 206.
OST_FIM	The name of the snapshot method that is selected in a policy configured for snapshot replication using Replication Director. The name represents OpenStorage Frozen Image Method. Refer to the NetBackup Replication Director Solutions Guide for more details.

Disk array configuration tasks

Note the following tasks.

Configuration tasks for the array administrator

Before you configure a backup, your array administrator or network administrator must do several tasks. Assistance with these tasks may be found in your array documentation or Windows documentation.

The array administrator must do the following tasks:

- Install the disk array and its software (such as Web) interface, including appropriate licenses.
- Install supported HBAs on the NetBackup primary client and alternate clients.
- If not already done, zone the client HBAs through the Fibre Channel switch, so the array is visible to the primary client and any alternate clients.
- Register your NetBackup primary client and any alternate clients with the array.
- Install NetBackup and array vendor snapshot management software on the NetBackup primary client and any alternate clients.
- Configure source and snapshot devices on the array (such as LUNs).
- For Microsoft Windows primary clients and any alternate clients, install the appropriate VSS provider for your array.

Configuration tasks for the NetBackup administrator

The NetBackup administrator must do the following tasks:

- If necessary, configure target devices on the array.
See the topic for your array.
See “Initial configuration of certain arrays” on page 158.
- For certain arrays, configure NetBackup disk array host credentials, which are required to access the array snapshot management software.
See the topic for your array.

EMC CLARiiON	See “Configuring NetBackup to access the CLARiiON array” on page 166.
EMC Symmetrix	See “About configuring NetBackup to access the Symmetrix array” on page 180.
HP EVA	See “Configuring NetBackup to access the EVA array” on page 188.
IBM DS6000 and DS8000	See “Configuring NetBackup to access the IBM DS6000 or DS8000 array” on page 190.
IBM DS4000	See “Configuring NetBackup to access the IBM DS4000 array” on page 199.
Hitachi SMS/WMS/AMS, USP/NSC, USP-V/VM	See “About communication between NetBackup and the Hitachi array” on page 202.
HP-XP	See “Determining if the HP-XP command devices are visible” on page 207.

- Create a NetBackup Snapshot Client policy for the array.
See the topic on the NetBackup policy for your array.

EMC CLARiiON	See “Configuring a NetBackup policy for a CLARiiON array method” on page 175.
EMC Symmetrix	See “Configuring a policy for EMC_TimeFinder methods” on page 182.
HP EVA	See “Configuring a NetBackup policy for an HP EVA array method” on page 189.
IBM DS6000 and DS8000	See “Configuring a NetBackup policy for IBM_DiskStorage_FlashCopy” on page 195.

IBM DS4000

See “Configuring a NetBackup policy for IBM_StorageManager_FlashCopy” on page 200.

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

See “Configuring a NetBackup policy for Hitachi_ShadowImage or Hitachi_CopyOnWrite” on page 204.

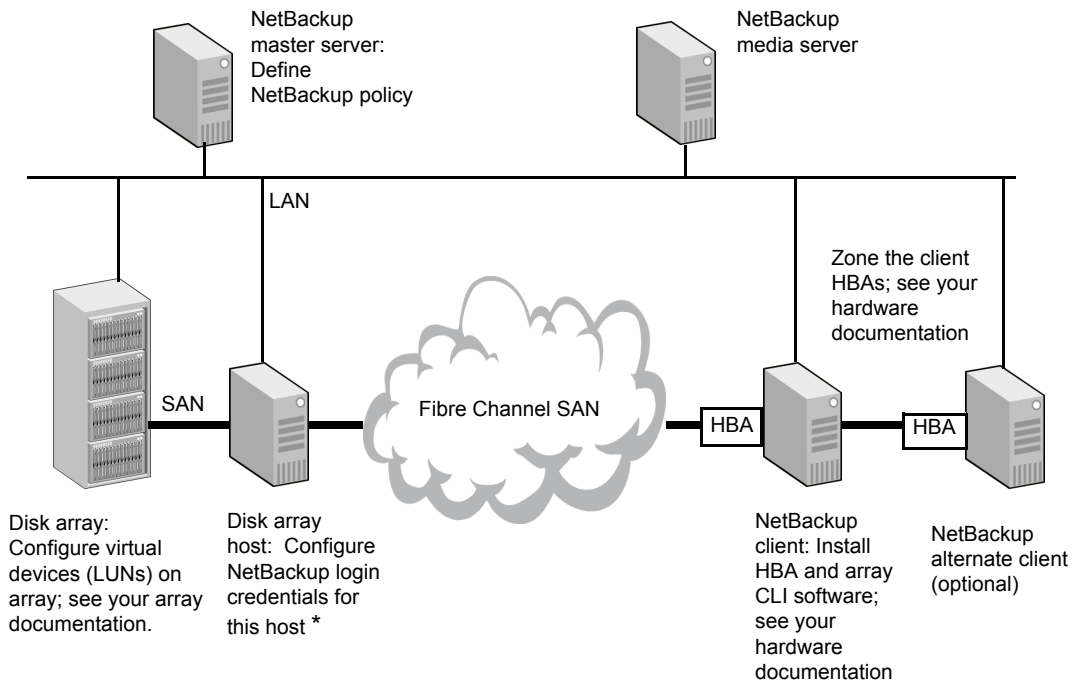
HP-XP

See “Configuring a NetBackup policy for HP_XP_BusinessCopy and HP_XP_Snapshot” on page 208.

Disk array configuration tasks diagram

The following diagram shows the major components and tasks that are required for a snapshot backup of a disk array. Some of these tasks must be performed by the array administrator.

Figure 9-1 Components for backup of disk array on Fibre Channel



* Some arrays do not have a separate front-end host; NetBackup credentials must be defined for the array itself. The array validates the NetBackup credentials.

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.

About Solaris `sd.conf` file

The `/kernel/drv/sd.conf` file must have sufficient entries to allow for the dynamic import of snapshot devices. Snapshot devices are created when the backup starts, and must be made visible on the NetBackup client without restarting the operating system. Entries should be added to the `sd.conf` file for the persistent target numbers that were configured for the HBA.

Note: The `sd.conf` file does not have to be modified if you use Leadville drivers.

Veritas recommends that you add LUNs 0-15 for all disk array targets on which snapshots are to be created. This creates 16 host-side LUNs on each target that can be used for importing the snapshots (clones, mirrors, and copy-on-write snapshots) required for backups. If 16 host-side LUNs are not enough for a particular disk array target, add more LUNs for that target. Note that snapshots are imported to a NetBackup client in sequential order starting with the lowest unused host-side LUN number. The host-side LUN number pool is managed on the disk array. The disk array cannot determine which host-side LUN numbers have been configured in `sd.conf`. The array can only determine which host-side LUN number it has not yet assigned to the host. If the array adds a device at a host-side LUN number that has not been configured in `sd.conf`, that device is not visible on the host. Also, if alternate client backups are being used, be sure to properly configure `sd.conf` on the alternate client.

You must restart after modifying `sd.conf`.

Symmetrix arrays pre-assign host-side LUN numbers (that is, the LUN numbers are not set at the time the device is imported). These pre-selected LUN numbers must be entered into `sd.conf` for the Symmetrix target number.

Note: If you use EMC Control Center interface (ECC) to determine Symmetrix host-side LUN numbers, note that ECC shows host-side LUN numbers in hexadecimal format. Since the LUN entries in `sd.conf` must be in decimal format, convert the hexadecimal value to decimal before adding it to `sd.conf`.

If the Symmetrix array was persistently bound at target 5, and the host-side LUN numbers of the Symmetrix devices are 65, 66, 67, then the following entries should be added to `sd.conf`.

```
name="sd" class="scsi" target=5 lun=65;
name="sd" class="scsi" target=5 lun=66;
name="sd" class="scsi" target=5 lun=67;
```

Solaris sd.conf file (Hitachi arrays only)

Before running backups, you must configure a sufficient number of static devices (array LUNs) in the `/kernel/drv/sd.conf` file on the client (and any alternate client) to accommodate the number of snapshot devices that backups require.

Note: Hitachi arrays do not support dynamic import of snapshot devices.

Linux modprobe.conf file

The `/etc/modprobe.conf` file must be set to allow the Fibre Channel HBA driver to scan for LUNs greater than 0. Make sure the following line (or something similar) appears in the `modprobe.conf` file:

```
options scsi_mod max_luns=255
```

If the line is not present, add it to the `modprobe.conf` and enter the following:

```
#mv /boot/initrd-linux_kernel_version.img  
/boot/initrd-linux_kernel_version.img.bak  
  
#mkinitrd -v /boot/initrd-linux_kernel_version.img  
linux_kernel_version
```

where the `linux_kernel_version` is the value that is returned from `uname -r` (for example, 2.6.9-34.ELsmp).

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
```

- Windows

```
\Program Files\Common Files\Veritas  
Shared\VxFI\4\Bin\nbfirescan.exe
```

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

Note the following regarding CLARiiON:

- If there are LUNs in the client's CLARiiON storage group, the LUNs are included in the output.
- If there are no LUNs visible but the array is zoned to allow the host to see it, the output includes the entry DGC LUNZ. This entry is a special LUN that the CLARiiON uses for communication between the client and the array. The LUNZ entry is replaced by another disk entry as soon as one is put in the storage group which has been presented to the client.

Example Solaris output, followed by a description:

DevicePath	Vendor	Product ID	EnclosureId	DeviceId	[Ctl,Bus,Tgt,Lun]
/dev/rdisk/c3t4d57s2	EMC	SYMMETRIX	000187910258	013C	[00,00,00,00]
/dev/rdisk/c3t6d10s2	HP	HSV200	5000-1FE1-5007-0020		
6005-08B4-0010-5F49-0000-5000-408F-0000					[00,00,00,00]

Note: The last line of output is wrapped.

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 Enclosure ID/DeviceID pair constitutes a client 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.

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]
```

- Most of the output lines are wrapped.
- DGC designates a CLARiiON device.

About VSS configuration (Windows)

On Windows clients, VSS is used during backups to create snapshots on disk arrays. Certain preconfiguration steps may be required as described in these sections.

Note on NetBackup array credentials

For certain disk arrays, you must supply NetBackup with logon credentials so that it can access the array. See the appropriate section in this chapter to configure NetBackup to access the array.

Note: 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.

Initial configuration of certain arrays

Certain arrays require some initial configuration to use the VSS method for a backup.

CLARiiON	No VSS-related preconfiguration is needed, except installation of the required array software and the VSS provider. See “Veritas support for VSS Snapshot and EMC CLARiiON” on page 163.
HP EVA	No VSS-related preconfiguration is needed, except installation of the required array software and the VSS provider. For any other configuration requirements, see the section in this chapter for your array.

Symmetrix	<p>You must associate the source device in the array with the target device(s) that are to be used for the differential (copy-on-write) or plex-based (clone or mirror) backup.</p> <p>Note: For Symmetrix arrays, NetBackup supports VSS with differential (copy-on-write) backup but not plex-based (clone or mirror) backup.</p>
EMC TimeFinder Snap	<p>See “Creating EMC disk groups for VSS differential snapshots that use EMC TimeFinder Snap” on page 159.</p> <p>See “Verifying that VSS can make a snapshot” on page 160.</p> <p>See “Testing the EMC TimeFinder Snap backup” on page 162.</p>

Creating EMC disk groups for VSS differential snapshots that use EMC TimeFinder Snap

Use the following procedure.

To create EMC disk groups for VSS differential snapshots that use EMC TimeFinder Snap

- 1 Create a disk group to contain any number of primary and secondary disks.

```
symdg create nbfim_test
```

Creates a disk group named nbfim_test.

- 2 Add primary disks to the disk group.

```
symld -g nbfim_test add dev 02A
```

Adds a primary disk 02A to disk group nbfim_test.

- 3 Add a VDEV disk to the disk group.

```
symld -g nbfim_test add dev 08C -vdev
```

Adds the VDEV disk 08C to the nbfim_test disk group.

When these commands are successfully entered, NetBackup can make snapshot backups with primary device 02A and its target VDEV 08C.

Verifying that VSS can make a snapshot

Before you finish preconfiguration or run a backup with the Windows VSS method, verify that the Volume Shadow Copy Service can make and delete a snapshot on the array.

The command that is used in the following procedure is available with Microsoft Volume Shadow Copy SDK 7.2.

On the NetBackup client, you can use the commands in the following procedure.

Note: If these commands do not succeed, consult your Windows documentation or your array administrator.

To verify that VSS can make or delete a snapshot

1 To create a snapshot, enter one of the following commands .

- For a differential (copy-on-write) snapshot:

```
vshadow.exe -ad -nw -p source_drive
```

- For a plex (clone or mirror) snapshot:


```
vshadow.exe -ap -nw -p source_drive
```

- 2** To display information on all existing snapshots on the client, enter the following command:

```
vshadow.exe -q
```

Example output:

```
VSHADOW.EXE 2.2 - Volume Shadow Copy sample client
Copyright ©) 2005 Microsoft Corporation. All rights reserved.
(Option: Query all shadow copies)
- Setting the VSS context to: 0xffffffff
```

```
Querying all shadow copies in the system ...
```

```
* SNAPSHOT ID = {ae445cc3-e508-4052-b0f6-a5f02cf85f1e} ...
  - Shadow copy Set: {6665f5f7-6468-4a22-bd73-29ef8a30a760}
  - Original count of shadow copies = 1
  - Original Volume name:
    \\?\Volume{0db3bc15-53b1-4d63-94dc-7c7d28b172cb}\ [K:\]

  - Creation Time: 4/10/2007 2:02:13 PM
  - Shadow copy device name:
    \\?\Volume{55226978-3131-4a12-8246-97ace27cf976}
  - Originating machine: oil.fun.com
  - Service machine: oil.fun.com
  - Not Exposed
  - Provider id: {21e5ab69-9685-4664-a5b2-4ca42bddb153}
  - Attributes: No_Auto_Release Persistent Hardware
No_Writers
Plex
```

- 3** To delete snapshots, do the following:

- Look in the `vshadow.exe -q` command output for the Shadow copy Set ID of a snapshot that you created above.
- Enter the following command to delete a particular snapshot:

```
vshadow.exe -dx = {Shadow_copy_set_ID}
```

- Enter the following command to delete all snapshots on the client:

```
vshadow.exe -da
```

Testing the EMC TimeFinder Snap backup

After you complete the Symmetrix configuration steps as described in this chapter and run a backup, you can verify that the snapshot was successful.

To test the EMC TimeFinder Snap backup

- ◆ Enter the following command to verify that the snapshot was successful.

```
symsnap -g nbfim_test query -multi
```

The `-multi` option shows multiple snapshots.

For example:

```
C:\Program Files\EMC\SYMCLI\bin\symsnap -g dmx-5D query -multi
```

where `dmx-5D` is the name of the Symmetrix device group.

Sample output:

```
Device Group (DG) Name: dmx-5D
DG's Type           : REGULAR
DG's Symmetrix ID    : 000187910258

Source Device      Target Device      State      Copy
-----
Protected          Changed
Logical  Sym  Tracks  Logical  Sym  G      Tracks      SRC <=> TGT  (%)
-----
DEV001      005D  273108  VDEV001  01A8      X      CopyOnWrite  0
              276208  VDEV002  01A9      X      2 CopyOnWrite  0

Total
Track(s)      552417      3
MB(s)         17263.0      0.1
```

If the `SRC <=> TGT` value reads `CopyOnWrite`, the snapshot was created successfully.

About EMC CLARiiON arrays

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

EMC CLARiiON software requirements for UNIX

Table 9-2 shows the required EMC software.

For software versions used in test configurations, see the *NetBackup Snapshot Client Configuration* document at:

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

Table 9-2 Software that is required for EMC CLARiiON

Software	Where to install
Navisphere Secure CLI	NetBackup clients
Navisphere Agent	NetBackup clients
CLARiiON SnapView software	EMC disk array
CLARiiON FLARE operating environment	EMC disk array

Veritas support for VSS Snapshot and EMC CLARiiON

Veritas has an open support policy for VSS Snapshot for NetBackup Snapshot Client. If a vendor supports a VSS provider for a Windows platform, Veritas provides support for local snapshot, alternate client, FlashBackup local snapshot, and FlashBackup alternate client methods. To use a CLARiiON disk array with VSS, contact EMC Corporation for the required software and versions. EMC supplies this software as a bundle, to ensure that the software components are at the right level and function correctly.

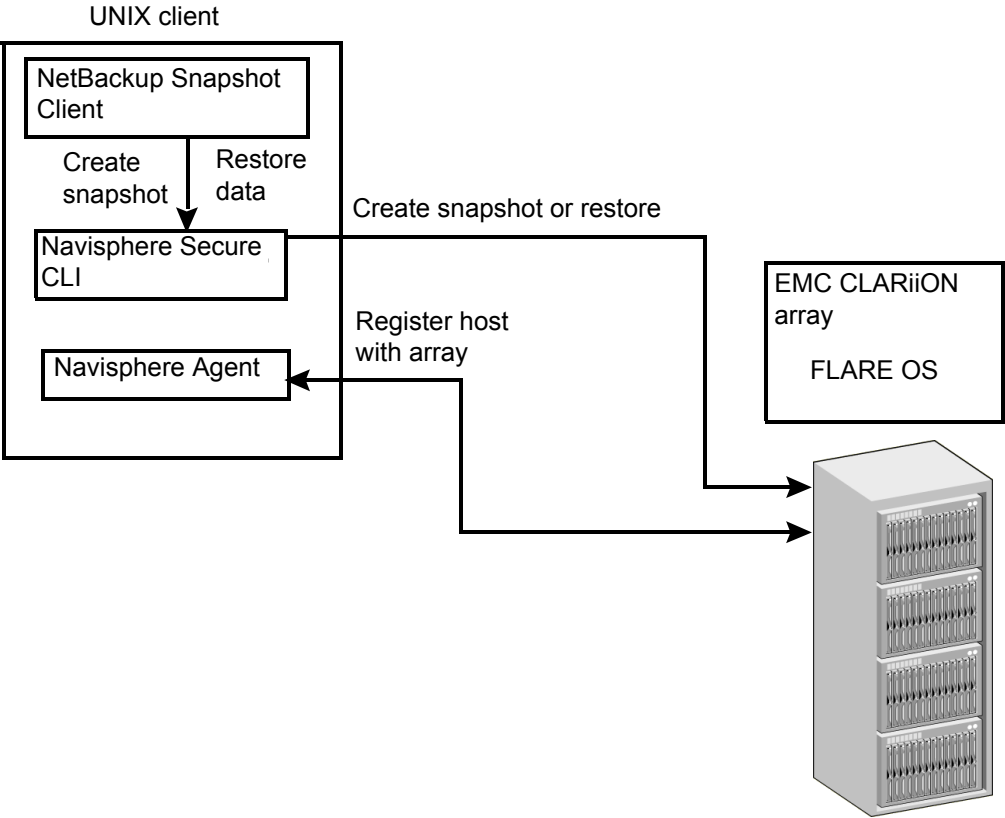
Note that the open support policy for VSS providers is not applicable to Instant Recovery. To use VSS along with the NetBackup Instant Recovery feature, refer to the NetBackup 7.x Snapshot Client Compatibility List for the components that NetBackup supports for Instant Recovery with the array. The compatibility list is available at the following URL:

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

Diagram of installed software for EMC CLARiiON

Figure 9-2 shows the software components on the NetBackup clients and the CLARiiON array for UNIX, and indicates the control function of each.

Figure 9-2 Software on NetBackup clients and CLARiiON array



Verifying connectivity from client to array

NetBackup communicates with the array by means of the EMC Navisphere Secure CLI. To verify that the CLI is installed and that NetBackup can communicate with the EMC CLARiiON array, enter the following command on each NetBackup client. Note that if a Navisphere security file has been created, the user name, password, and scope do not have to be supplied with the command.

To verify connectivity from client to array

1 Enter the following:

■ On UNIX:

```
/opt/Navisphere/bin/naviseccli -h CLARiiON_hostname -user  
array_admin_username -password array_admin_password -scope 0  
getagent
```

■ On Windows:

```
Program Files\EMC\Navisphere CLI\naviseccli -h CLARiiON_hostname  
-user array_admin_username -password password -scope 0 getagent
```

Sample output from this command:

```
Agent Rev:          6.19.1 (2.6)  
Name:              K10  
Desc:  
Node:              A-APM041147024  
Physical Node:     K10  
Signature:         1218092  
Peer Signature:    1099522  
Revision:          2.19.500.5.027  
SCSI Id:           0  
Model:             CX500  
Model Type:        Rackmount  
Prom Rev:          3.60.00  
SP Memory:         2048  
Serial No:         APM041147024  
SP Identifier:     A  
Cabinet:           DPE2
```

If the command fails, you must address the problem before you do any further array configuration.

This problem could be due to the following:

- The Navisphere Secure CLI or other array interface is not installed.
- The installed Navisphere Secure CLI is not a supported version.
See Table 9-2 on page 163.
- The array credentials are invalid.

Note: On AIX or some UNIX host the Snapshot creation can fail for EMC_CLARiiON array, if the Navisphere Secure CLI location entries are incorrect in the `/usr/opensv/lib/vxfi/configfiles/emccлариionfi.conf` file.

For example, on the AIX host `naviseccli` is found at the following location `/usr/lpp/NAVICLI/naviseccli`. Verify the correct `naviseccli` path and add the following file path and name entries to the `/usr/opensv/lib/vxfi/configfiles/emccлариionfi.conf` file.

- `FILEPATH_NAVISEC_EXE="filepath"`
- `FILENAME_NAVISEC_EXE="filename"`

- 2 For further troubleshooting information, consult your array documentation.

About resolving host names on the network

All NetBackup clients to be used with Snapshot Client must be resolvable by means of your network naming service (such as DNS). NetBackup uses the network naming service to look up the client's host name, its fully qualified domain name, and IP address. NetBackup then looks for any of those three in a CLARiiON storage group when attempting to do LUN masking.

Configuring NetBackup to access the CLARiiON array

You must supply logon credentials that allow the NetBackup client to access each storage processor on each EMC CLARiiON array.

IMPORTANT: For single path configurations, you must add credentials for the CLARiiON storage processor that owns the source LUNs and any clone LUNs. For multipath configurations, you must add credentials for both storage processor A and storage processor B. The storage processor that owns the LUN may be changed by the multipath software. (For single path configurations, all CLARiiON LUNs must be owned by only one storage processor.)

To configure NetBackup to access the CLARiiON array

- 1** In the NetBackup Administration Console, click the **Media and Device Management > Credentials > Disk Array Hosts** node in the NetBackup Administration Console.
- 2** Right-click in the **Disk Array Hosts** pane and select **New Disk Array Host**.
- 3** Enter the host name of the EMC CLARiiON array.
- 4** Select EMC CLARiiON in the **Disk Array Host Type** pull-down menu.
- 5** Enter the user name and password for the storage processor.
- 6** Clear (uncheck) the **Connect using port number** checkbox.

About using a local scope CLARiiON administrator account

By default, NetBackup uses a global scope CLARiiON administrator account, not a local scope account. To specify a local scope CLARiiON administrator account, you must use the Navisphere security file to supply the account information (see the following note, however). The Navisphere security file contains a user name, password, and scope for the CLARiiON administrator. The CLARiiON snapshot provider uses the Navisphere security file if the file exists. To create the security file, see the EMC Navisphere documentation.

Note: You must also enter credentials by means of the **Disk Array Hosts** dialog box in the NetBackup Administration Console. The disk array host name is not provided in the Navisphere security file.

Adding clients to a CLARiiON storage group

All NetBackup primary clients and any alternate clients must exist in a CLARiiON storage group. NetBackup does not automatically add clients to a CLARiiON storage group.

Warning: Veritas strongly recommends that every NetBackup client be given its own CLARiiON storage group on the array. Data corruption could result if more than one client (host) exists in a single storage group. If it is necessary to have multiple hosts in a single storage group, you must make certain that only one host in the storage group is actually using the device at any given time. (Only one host should mount the disk.) A Windows host may actually write to a LUN masked device even if the device is not mounted. Therefore, a Windows host should always be in its own storage group.

To add clients to a CLARiiON storage group

- 1 Register your clients (hosts) with the array.
 - 2 Create storage groups.
 - 3 Add the NetBackup primary and any alternate clients to the storage groups.
- For more detail, see your array administrator or array documentation.

Configuring for EMC_CLARiiON_SnapView_Clone

The following must be completed before NetBackup can use the clone capabilities of the CLARiiON array. The following steps are described in the topics that follow.

Table 9-3 Configuration process for EMC_CLARiiON_SnapView_Clone

Step	Action	Related topic
Step 1	Array administrator creates clone private LUNs.	See “Creating a clone private LUN with the EMC Navisphere Web interface” on page 169.
Step 2	Array administrator creates a clone group and selects a LUN as source.	See “Creating a clone group and select a LUN as source” on page 169.
Step 3	Array administrator adds clone LUNs to the clone group.	See “Adding clone LUNs to the clone group” on page 170.
Step 4	Array administrator supplies source and target devices.	See “Obtaining the device identifier for each source and clone LUN” on page 172. See “About configuration for EMC_CLARiiON_SnapView_Snapshot” on page 173. See “Configuring a reserved LUN pool for the storage processors” on page 174.
Step 5	NetBackup administrator configures a NetBackup policy for the array, using the device that are identifiers supplied by the array administrator.	See “Configuring a NetBackup policy for a CLARiiON array method” on page 175.

Note: For Windows clients and the VSS method, you must synchronize the clone with its source.

Note: These steps are separate from those taken by NetBackup to create the backup. When the backup begins, NetBackup synchronizes the clones with the source (if necessary) and splits (fractures) the clones to make them available for the backup.

For more information on the EMC array terminology in this section, see your EMC CLARiiON documentation.

Creating a clone private LUN with the EMC Navisphere Web interface

You must configure a clone private LUN for each CLARiiON storage processor that owns a clone source LUN. Clone private LUNs store the portions of the client's data that incoming write requests change while the clone is in use. Clone private LUNs are used while a clone LUN is fractured and when a synchronization occurs.

A clone private LUN can be any bound LUN that is at least 250,000 blocks in size.

To create a clone private LUN with the EMC Navisphere Web interface

- 1 Right-click the array name.
- 2 Right-click the **Snapview** node and select **Clone Feature Properties**.
- 3 Choose the LUNs you want to label as Clone Private LUNs.

Choose a clone private LUN for each storage processor that contains clone source LUNs. (You must know which storage processor owns a given LUN.) Only one clone private LUN is required per storage processor. You can add more clone private LUNs later if more space is needed.

Creating a clone group and select a LUN as source

Use the following procedure to create a clone group and select a LUN as source.

To create a clone group and select a LUN as source

- 1 In the EMC Navisphere Web interface, right-click the **Snapview** node and select **Create Clone Group**.

Name	ID	Capacity	Drive Type
LUN 7	7	33.000GB	Fibre Channel
LUN 8	8	0.250GB	Fibre Channel
LUN 14	14	0.125GB	Fibre Channel
LUN 17	17	0.098GB	Fibre Channel
LUN 24	24	0.250GB	Fibre Channel
LUN 38	38	0.500GB	Fibre Channel
LUN 42	42	1.205GB	Fibre Channel
LUN 50	50	2.000GB	Fibre Channel
LUN 61	61	500.000GB	Fibre Channel
LUN 66	66	250.000GB	Fibre Channel
LUN 67	67	250.000GB	Fibre Channel

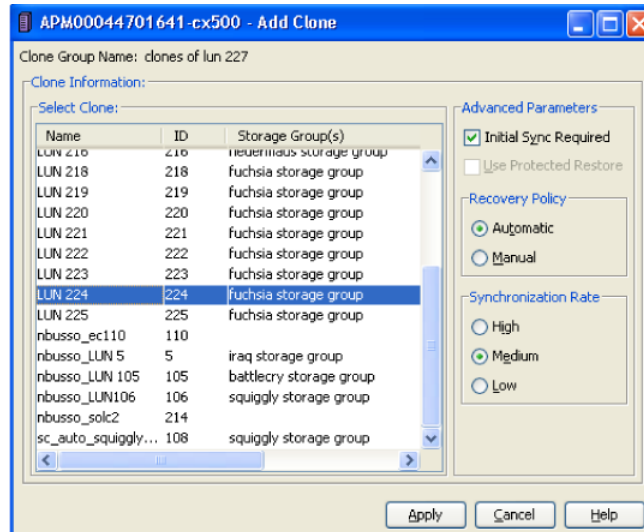
- 2 Enter a name for the clone group.
- 3 Choose the source LUN to be cloned (the source LUN to associate with this group). This LUN contains the data that is copied to the clones.

Adding clone LUNs to the clone group

Use the following procedure.

To add clone LUNs to the clone group

- 1 In the EMC Navisphere Web interface, on the **SnapView > Clones** node, right-click the name of the clone group and select **Add Clone**.



- 2 Select a LUN from the list to be added to the clone group.
 Accept the defaults for **Advanced Parameters**, **Recovery Policy**, and **Synchronization Rate**.
 - Do not choose a LUN that contains valuable data. Any data on the LUN is overwritten by this procedure.
 - The number of clones to add depends on how many point-in-time copies of the source LUN you want to retain at one time. The limit is 8 clones per source LUN.
 This value should match the **Maximum Snapshots** setting in the EMC_CLARiON_Snapview_Clone configuration parameters in the **Snapshot Options** dialog in a NetBackup policy.
 See "Maximum Snapshots parameter" on page 101.
- 3 When you click **Apply**, Navisphere begins to copy data from the source LUN to the LUN you have selected, creating a clone LUN.
 Any previous data on the clone LUN is lost.

Obtaining the device identifier for each source and clone LUN

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 follows.

To obtain the device identifier for each source and clone LUN

- 1 Enter the following command on the NetBackup client:

```
/opt/Navisphere/bin/naviseccli -address CLARiON_hostname -user
array_admin_username -password password -scope 0 getlun
lun_number -uid
```

- 2 Note the exact UID string that this command returns. This UID is the unique ID of the LUN.

For example, to obtain the unique ID of LUN 67, enter:

```
/opt/Navisphere/bin/naviseccli -address CLARiON_hostname -user
array_admin_username -password password -scope 0 getlun 67
-uid
```

Example output:

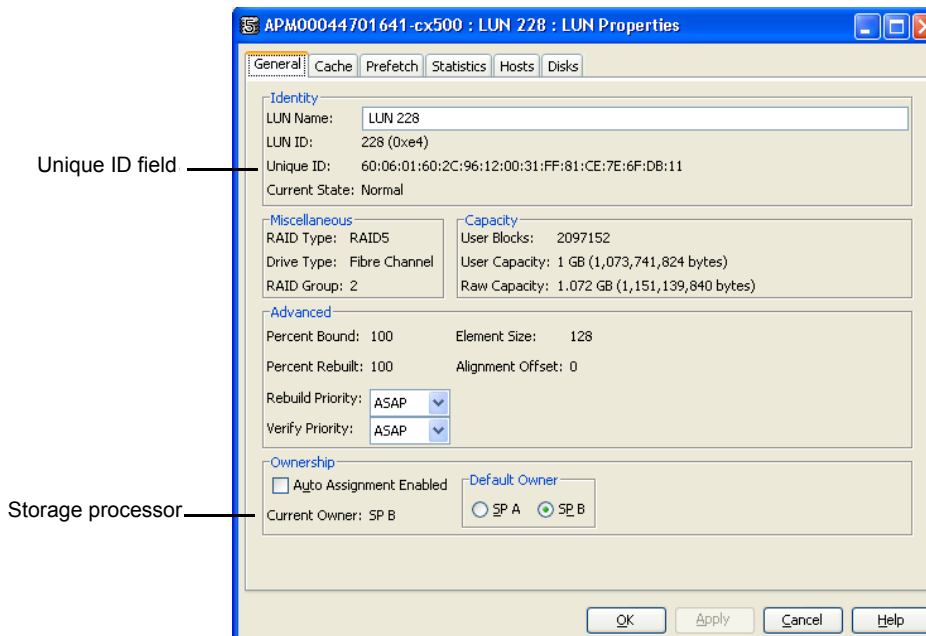
```
UID: 60:06:01:60:C8:26:12:00:4F:AE:30:13:C4:11:DB:11
```

- 3 To obtain the number of the LUN to use on the naviseccli command, find the clone group and examine the LUN list.
- 4 Copy the unique ID into the NetBackup policy, as follows:
 - If the LUN specified on the naviseccli command is the source LUN for the clone group, copy the unique ID into the **Source Device** field of the **Add Snapshot Resource** dialog box of the NetBackup policy. Help is available for that dialog box.
 See "Configuring a policy using EMC_CLARiON_Snapview_Clone method" on page 175.
 - If the LUN specified on the naviseccli command is a clone LUN, copy the unique ID into the **Snapshot Device(s)** field.

To use the Navisphere Web interface

- 1 To find the unique ID of a LUN, right-click the LUN in the Navisphere Web interface and select **Properties**.
- 2 See the **Unique ID** field.

As of this writing, you cannot copy and paste the Unique ID from the Navisphere web interface. Use the `naviseccli` command to copy and paste the ID.



About configuration for EMC_CLARiiON_SnapView_Snapshot

You must configure a reserved LUN pool for each storage processor that owns a source LUN for which you want to create a SnapView snapshot. (Starting at FLARE version x.24, one combined reserved LUN pool exists for both storage processors.) The reserved LUN pool stores the portions of the client's data that incoming write requests change while the snapshot is active.

Note the following before you configure a reserved LUN pool.

- The reserved LUN pool must contain at least one LUN for each source LUN that is the subject of (participates in) the snapshot. Any available LUN can be added to the reserved LUN pool if the LUN is owned by the storage processor that owns the source LUN.

- LUNs in the reserved LUN pool are private LUNs, which cannot belong to a storage group. The storage processor manages its reserved LUN pool and automatically assigns one or more private LUNs to a source LUN. This assignment is based on how much snapshot activity takes place in the source LUN. This activity can result from one busy snapshot or multiple snapshots.
- While the snapshot is active, client write activity on the source consumes more space in the reserved LUN pool. Adding more LUNs to the reserved LUN pool increases the size of the reserved LUN pool. The storage processor automatically uses a LUN if one is needed.
- All snapshots share the reserved LUN pool. If two snapshots are active on two different source LUNs, the reserved LUN pool must contain at least two private LUNs. If both snapshots are of the same source LUN, the snapshots share the same private LUN (or LUNs) in the reserved LUN pool.

Configuring a reserved LUN pool for the storage processors

In the EMC Navisphere Web interface, do the following.

To configure a reserved LUN pool for the storage processors

- 1 Under the node for your array, right-click **Reserved LUN Pool** and choose **Configure**.
- 2 In the **Configure Reserved LUN Pool** dialog box, under **Available LUNS**, select the LUNs to be added to the reserved LUN pool.

The reserved LUN pool must contain at least one LUN for each source LUN.
- 3 For FLARE versions before 2.24, click the **Add to SPA LUN Pool** for storage processor A or **Add to SPB LUN Pool** for storage processor B.

Check the properties of the source LUN to determine the storage processor that owns it.
- 4 Click **Ok** or **Apply**.

Each reserved LUN pool node shows how much space is free in the pool.

The amount of space that the reserve LUN pool requires depends on how many unique blocks change in the source LUN during the life of the snapshot. If little space remains in an existing LUN pool, add new LUNs to the LUN pool.

Generally, the more snapshots that are created, the more space that is required in the reserved LUN pool. If the LUN pool space is used up, snapshots attached to backups may become inactive and future snapshot backups fail.

Configuring a NetBackup policy for a CLARiiON array method

General help for setting up a NetBackup policy is available.

See “Configuring a Snapshot Client policy” on page 52.

To configure a NetBackup policy for a CLARiiON array method

- 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 a CLARiiON method. To configure the method, see one of the following topics.

Configuring a policy using EMC_CLARiiON_Snapview_Clone method

In the **Snapshot Options** dialog box of the policy, you must specify the following information in the **Snapshot Resources** pane. NetBackup uses this information to correctly rotate through the clone LUNs when performing a SnapView 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.

To configure a policy using EMC CLARiiON Snapview Clone method

- 1 Obtain the following source and clone LUN information:
 - Array serial number (in the EMC Navisphere Web interface, right-click the array name and select **Properties**).
 - Unique ID for the source LUN containing the primary data.
 - Unique ID for the target LUN(s) used for snapshots.

See “Obtaining the device identifier for each source and clone LUN” on page 172.

Or see your array administrator.

- 2 Copy and paste (or type) these values into the appropriate field of the NetBackup policy’s **Add Snapshot Resource** dialog box, as explained below.

The policy Backup Selections specify a set of 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, the procedure is as follows.

See “Configuring the EMC_CLARiiON_Snapview_Clone method” on page 176.

Configuring the EMC_CLARiiON_Snapview_Clone method

Use the following procedure to configure the EMC_CLARiiON_Snapview_Clone method.

To configure the EMC_CLARiiON_Snapview_Clone method

- 1 After you select **EMC_CLARiiON_Snapview_Clone** on the **Snapshot Options** dialog box, click **Add**.
- 2 In the **Add Snapshot Resource** dialog box, enter the array's serial number in the **Array Serial #** field.
- 3 Enter the unique ID for the source LUN in the **Source Device** field.
- 4 Enter the unique IDs for the clone LUNs in the **Snapshot Device(s)** field. To enter multiple IDs, place a semicolon between them.

For Instant Recovery backups, the Snapshot Device(s) entries determine where and in what order the snapshots are retained.

See "Snapshot Resources pane" on page 100.

EMC_CLARiiON_Snapview_Snapshot method

In the **Snapshot Options** dialog box of the policy, you can set the **Maximum snapshots (Instant Recovery only)** parameter for the **EMC_CLARiiON_Snapview_Snapshot** method. The maximum value is 8.

See "Maximum Snapshots parameter" on page 101.

Common CLARiiON array configuration problems

Note the following regarding the CLARiiON array:

Do not use the EMC_CLARiiON_Snapview_Clone method and the EMC_CLARiiON_Snapview_Snapshot method to back up the same source LUN. If you attempt a rollback restore from a EMC_CLARiiON_Snapview_Snapshot snapshot, the restore fails if a clone associated with the source LUN is currently synchronized.

About EMC Symmetrix arrays

The following topics include background information and configuration tasks for snapshot backups using EMC Symmetrix arrays. These tasks must be completed before you run a backup.

EMC Symmetrix DMX software requirements

Table 9-4 shows the required EMC software.

Table 9-4 Software that is required for EMC Symmetrix

Software	Where to install	Versions
EMC Solutions Enabler	NetBackup clients	For versions used in test configurations, see <i>Veritas NetBackup Snapshot Client Configuration</i> , at: http://www.veritas.com/docs/000081320
Symmetrix Solutions Enabler license	NetBackup clients	For versions used in test configurations, see <i>Veritas NetBackup Snapshot Client Configuration</i> , at: http://www.veritas.com/docs/000081320

Clone emulation flag can cause snapshots to fail

Clone emulation flag can cause snapshots to fail:

- The operating system is Windows.
- An EMC DMX array is used.
- The BCV relationships were created with the clone emulation flag set to `TRUE`.

Given this situation, snapshots may fail if the `EnableCloneEmulation Windows` registry entry is not set to `TRUE`. For snapshots to succeed, set the entry to `TRUE`.

This registry entry is found in the following

location: `HKEY_LOCAL_MACHINE\Software\EMC\ShadowCopy`

Registry details:

- **Name:** `EnableCloneEmulation`
- **Type:** `REG_SZ`

Possible values include:

- `TRUE`: Enables Symmetrix array clone emulation.
- `FALSE`: Disables Symmetrix array clone emulation.

EMC snapshot operation fails

Snapshot fails if the `EnforceStrictBCVPolicy` registry entry is not set to `TRUE`. For snapshots to succeed, set the entry to `TRUE`. This registry entry is found at the following location:

HKEY_LOCAL_MACHINE\Software\EMC\ShadowCopy

Registry details:

- **Name:** EnforceStrictBCVPolicy
- **Type:** REG_SZ

Possible values include:

- **TRUE:** Indicates that EMC VSS Provider enforces a strict BCV rotation policy, where a BCV should only be used if it is not currently part of a snapshot.
- **FALSE:** Indicates that EMC VSS Provider does not enforce a BCV rotation policy, leaving enforcement to the VSS requestor.

Support for EMC Symmetrix with Volume Shadow Copy Service

To use an EMC Symmetrix disk array with Volume Shadow Copy Service, contact EMC Corporation for the required software and versions. EMC supplies this software as a bundle, to ensure that the software components are at the right level and function correctly.

Prerequisites for using EMC Symmetrix disk arrays

Note the following prerequisites before using EMC Symmetrix disk arrays:

- The array administrator must allocate source and target devices before you create any NetBackup policies. See the array administrator for the device identifiers.
- The array administrator must allocate a gatekeeper device and a VCMDB (Volume Configuration Management Database) for all of the NetBackup clients. Storage servers should also have the appropriate number of gatekeeper devices configured.

Note: Gatekeeper devices are not mandatory for UNIX clients.

A VCMDB is a virtual LUN database that keeps track of which LUNs the client can see. A gatekeeper is a small disk that the DMX uses to pass commands between the client and the array.

Configuring NetBackup clients to use EMC Symmetrix

Use the following procedure to configure NetBackup clients to use EMC Symmetrix disk arrays.

To configure NetBackup clients to use EMC Symmetrix disk arrays

- 1 On Windows hosts that will be used to monitor devices, install the EMC Solutions Enabler and the gatekeeper devices. In general, NetBackup production clients are not assigned gatekeepers and are monitored from a remote host.

The EMC Solutions Enabler is not needed on UNIX clients. Gatekeeper devices are not mandatory for UNIX clients.

- 2 Register your EMC license keys using the `symlmf` command.
- 3 Discover the HBAs in the NetBackup client, to allow the NetBackup client to perform LUN masking operations in the array.

To discover host HBAs, run the following Solutions Enabler SYMCLI command:

```
symmask.exe discover hba
```

Example output:

```
Symmetrix ID          : 000292603831
Device Masking Status : Success
```

Identifier	Type	User-generated Name
2100001b3212c04f	Fibre	2100001b3212c04f/2100001b3212c04f

```
Symmetrix ID          : 000492600276
Device Masking Status : Success
```

Identifier	Type	User-generated Name
2100001b3212c04f	Fibre	2100001b3212c04f/2100001b3212c04f

If no Symmetrix ID appears in the output, there is a connectivity problem.

If the command fails, you must address the problem before doing any further array configuration. This problem can be due to (but is not necessarily limited to) the following:

- The Solutions Enabler interface is not installed.
- Your NetBackup client may not be properly zoned to see the array.

For further troubleshooting information, consult your array documentation.

About configuring NetBackup to access the Symmetrix array

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

About configuration for EMC_TimeFinder_Mirror

The EMC_TimeFinder_Mirror method uses the Symmetrix Business Continuance Volume (BCV) to allow NetBackup to perform mirror backups. For this method, the Snapshot Client source devices are assumed to be STD devices and the target snapshot resources are assumed to be BCV devices. BCV devices are normally created by your EMC technical representative. The EMC_TimeFinder_Mirror method requires that the source device or standard device (STD) be fully synchronized (established) with the snapshot device (BCV). BCV devices must already exist before they can be synchronized and used in a NetBackup backup.

Fully synchronizing STD/BCV mirror pairs

Make sure that each target (BCV) disk is fully synchronized with its source (STD). One way to accomplish this synchronization is as follows.

To fully synchronize STD/BCV mirror pairs

- 1 Create a temporary file that contains only the source and target device IDs separated by a space. (Only one source-target pair can exist in a temporary file.)

For example, if the source (STD) device ID is 0050 and the target (BCV) device ID is 0060, the temporary file should contain the following:

```
0050 0060
```

- 2 Use the `symmir` command to fully establish the mirror pair.

```
symmir -sid 000187910258 establish -f temp_file -full
```

When the pair is synchronized, it can be used in a NetBackup policy. Synchronization can take a long time. For example, it may take between 5 and 10 minutes for an 8GB STD/BCV pair to enter the synchronized state.

3 Check the status of the mirror pair:

```
symmir -sid 000187910258 query -file temp_file
```

Make sure the *temp_file* name matches the *temp_file* name you used above.

4 In the output, look for *Synchronized* under the *State* column. When the pair enters the synchronized state, it is ready to be used for backup.

About configuration for EMC_TimeFinder_Clone

The EMC_TimeFinder_Clone method uses Symmetrix standard devices (STD) as the NetBackup Snapshot Client source devices and TimeFinder Clone copies as the snapshot devices. Any Symmetrix STD or BCV device can be used as an EMC_TimeFinder_Clone method snapshot resource.

This method does not require the establishment of a clone relationship between the primary and secondary before the backup. NetBackup automatically manages the primary-secondary clone relationship in the context of a Symmetrix copy session.

Verifying that the clone is complete before doing a point in time rollback

NetBackup initiates clone creation on the array. For large source devices, clone creation may take a long time to complete. The Symmetrix array, however, makes the new clone immediately available for normal restore. This availability allows individual files or directories to be restored as soon as the NetBackup Activity Monitor marks the job as finished, even if the clone has not been fully created.

Although you can restore individual files or directories before the clone is completely copied, you cannot perform a point in time rollback. Clone creation must be complete before a point in time rollback can succeed.

If a point in time rollback begins before the clone is finished, the restore fails without a failure notification in the Activity Monitor. In this case, the *bpfis* log contains the phrase *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.

To verify that the clone is complete before doing a point in time rollback

- 1 Create a temporary file that contains only the source and target device IDs separated by a space.

For example, if the source device ID is 0050 and the target (clone) device ID is 0060, the temporary file should contain the following:

```
0050 0060
```

- 2 Check the status of the clone with the `symclone` command. For example:

```
symclone -sid 58 query -file /tmp/0050_0060.txt
```

where 58 is a short version of a Symmetrix ID ending in 58, such as 000187910258, and `/tmp/0050_0060.txt` is the temporary file.

- 3 In the output, look for `Copied` under the `State` column. When the clone pair is in the copied state, it is ready for point-in-time rollback.

About configuration for EMC_TimeFinder_Snap

The EMC_TimeFinder_Snap method uses Symmetrix virtual devices (VDEV) to allow NetBackup to perform copy-on-write snapshot backups. This method does not require the establishment of a source-snapshot relationship between the primary and the VDEV devices before the backup. NetBackup automatically manages the source-snapshot relationship.

VDEV devices must already exist before they can be used in a NetBackup backup. Normally virtual devices are created by your EMC technical representative.

Configuring a policy for EMC_TimeFinder methods

Use the following instructions for the EMC TimeFinder array methods (EMC_TimeFinder_Snap, EMC_TimeFinder_Clone, EMC_TimeFinder_Mirror).

Help for setting up a NetBackup policy is available.

See “Configuring a Snapshot Client policy” on page 52.

To configure a policy for EMC_TimeFinder methods

- 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 a TimeFinder method.

Enter a value in the **Value** field for the **SAVE Device Pool**. You can also use the default pool.

When you use the default pool for all operations, the device space is quickly consumed and backups can fail. But when you use a separate pool to configure the policy, there is less chance of backup failure.

- 4 In the **Snapshot Options** dialog box of the policy, you must specify the following information:

- Symmetrix ID.
- Unique ID for the snapshot resource from which the TimeFinder_Mirror, Clone, or Snapshot is created.
- Unique ID for the target device(s) where the mirror, clone, or snapshot is created.

NetBackup uses this information to correctly rotate through the snapshot, clone, or mirrors when performing a backup. After all the snapshots, clones, or mirrors have been used in a backup, NetBackup automatically determines which is the oldest. NetBackup expires the oldest snapshot, clone, or mirror so that it can be reused for the current backup.

- 5 See your array administrator for these values.
- 6 Copy and paste (or type) these values into the appropriate field of the NetBackup policy's **Add Snapshot Resource** dialog box. The procedure is as follows.
- 7 After selecting an EMC_TimeFinder method on the **Snapshot Options** dialog box, click **Add**.
- 8 In the **Add Snapshot Resource** dialog box, enter the Symmetrix ID in the **Array Serial #** field.
- 9 Enter the unique ID for the source device in the **Source Device** field.
- 10 Enter the unique IDs for the target devices in the **Snapshot Device(s)** field.
To enter multiple IDs, place a semicolon between them. The ID should be a four-digit value.
 - For EMC_TimeFinder_Mirror, the target devices are BCV devices.
 - For EMC_TimeFinder_Snap, the target devices are VDEV devices.

- For EMC_TimeFinder_Clone, the target devices are the STD devices that were allocated to be used as clones.
- 11 Enter source and target device IDs exactly as they appear on the Symmetrix.
- For example, if device 4c appears as 004C, then enter it as 004C (case does not matter). The `symdev show` command can be used to determine how a device ID appears on Symmetrix. Refer to your SymCLI documentation for more information on this command.
- For Instant Recovery backups, the Snapshot Device(s) entries determine where and in what order the snapshots are retained.
- See “Snapshot Resources pane” on page 100.

About HP EVA arrays

The following sections describe configuration steps for supported HP arrays. These steps must be completed before you run a backup.

Prerequisites for working with HP EVA arrays

- Note the following prerequisites for the tasks in this section. See your array administrator for further assistance.
- Add your host (HBA) to the EVA array. All NetBackup primary clients and any alternate clients must be "added" to the EVA array. You can use the SSSU for HP StorageWorks Command View EVA utility or the StorageWorks Command View EVA Web interface. Refer to your HP EVA documentation for details.
- The EVA array requires that the NetBackup client's world-wide port name (not the world-wide node name) be used when you add the host entry to the array.

HP EVA software requirements for UNIX

Table 9-5 shows the required HP software.

Table 9-5 Software that is required for HP EVA

Software	Where to install	Versions
SSSU for HP StorageWorks Command View EVA (CLI)	NetBackup clients	For versions used in test configurations, see <i>Veritas NetBackup Snapshot Client Configuration</i> , at: http://www.veritas.com/docs/000081320

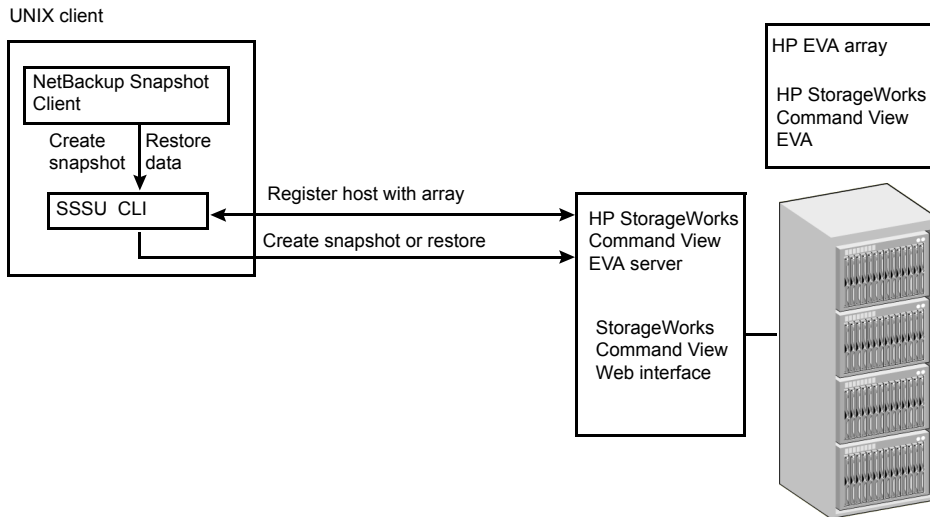
Table 9-5 Software that is required for HP EVA (*continued*)

Software	Where to install	Versions
HP StorageWorks Command View EVA Web interface	HP Command View EVA server	For versions used in test configurations, see <i>Veritas NetBackup Snapshot Client Configuration</i> , at: http://www.veritas.com/docs/000081320

Diagram of installed software for HP EVA

Figure 9-3 shows the software components on the NetBackup clients and the HP EVA Command View server for UNIX and indicates the control function of each.

Figure 9-3 Software components on NetBackup clients and HP EVA



Veritas support for VSS Snapshot and HP EVA

Veritas has an open support policy for VSS Snapshot for NetBackup Snapshot Client. If a vendor supports a VSS provider for a Windows platform, Veritas provides support for local snapshot, alternate client, FlashBackup local snapshot, and FlashBackup alternate client methods. To use an HP EVA disk array with VSS,

contact Hewlett Packard Enterprise for the required software and versions. HP supplies this software as a bundle, to ensure that the software components are at the right level and function correctly.

Note that the open support policy for VSS providers is not applicable to Instant Recovery. To use VSS along with the NetBackup Instant Recovery feature, refer to the NetBackup 7.x Snapshot Client Compatibility List for the components that NetBackup supports for Instant Recovery with the array. The compatibility list is available at the following URL:

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

Verifying connectivity from clients to array using SSSU 5.0

NetBackup communicates with the array by means of the SSSU for HP StorageWorks Command View EVA utility or the StorageWorks Command View EVA Web interface. To verify that the CLI is installed and that NetBackup can communicate with the array, use the following procedure on each NetBackup client.

To verify connectivity from clients to array using SSSU 5.0

- 1 Connect to the disk array host:

```
/opt/hp/sssu/sssu_sunos
```

Example output:

```
SSSU for HP StorageWorks Command View EVA 5.0 [v5.0.12]
```

- 2 Log on to the disk array host:

```
NoSystemSelected> select manager manager_host  
user=manager_username password=manager_password
```

- 3 Verify that you can see the EVA arrays that are managed by the host:

```
NoSystemSelected> ls cell
```

Example output:

```
Systems available on this Manager:  
HPEVA4000  
VRTS.EVA.ROS
```

See “Troubleshooting the SSSU procedure” on page 187.

To verify connectivity from clients to array using SSSU 6.0

1 Connect to the disk array host:

```
/opt/hp/sssu/sssu_sunos
```

Example output:

```
SSSU for HP StorageWorks Command View EVA  
Version: 6.0  
Build: 34
```

2 At the prompts, log on to the disk array host:

```
Manager:disk array host name  
Username:username  
Password:password
```

3 Verify that you can see the EVA arrays that are managed by the host:

```
NoSystemSelected> ls cell
```

Example output:

```
Systems available on this Manager:  
hpeva_2200nc07
```

See “Troubleshooting the SSSU procedure” on page 187.

Troubleshooting the SSSU procedure

If the SSSU command fails, you must address the problem before you do any further array configuration. This problem could be due to the following:

- The SSSU for HP StorageWorks Command View EVA utility is not installed.
- The SSSU for HP StorageWorks Command View EVA utility is not a supported version.
- The array manager credentials are invalid.

For further troubleshooting information, consult your EMC Storage Scripting System Reference manual.

Policy validation failure for the SSSU command

Policy validation fails when the SSSU CLI is not present at the expected path. The logs display the message

Unable to locate the CLI tool /opt/hp/sss/ at path sssu_hpx,

The cause of the error message is the CLI path, which is different from the default CLI path.

To fix the policy validation, add the following entry into the `hpevafi.conf` file:

```
[CLI_TOOL_INFO]
"FILEPATH"="<SSSU CLI path>"
"FILENAME"="<SSSU CLI tool file name>"
```

For example, for the HP platform the path would be:

```
/usr/opensv/lib/vxfi/configfiles/hpevafi.conf
```

The entry would be:

```
[CLI_TOOL_INFO]
"FILEPATH"="/opt/hp/sss/"
"FILENAME"="sss_hpx_parisc"
```

After you manually add these inputs to the `hpevafi.conf` file, the validation is successful.

Configuring NetBackup to access the EVA array

You must configure logon credentials on the NetBackup master server that allow the NetBackup client to access the array, as follows.

To configure NetBackup to access to the array

- 1 In the NetBackup Administration Console, click the **Media and Device Management > Credentials > Disk Array Hosts** node in the NetBackup Administration Console.
- 2 Right-click in the **Disk Array Hosts** pane and select **New Disk Array Host**.
- 3 Enter the host name through which the array management service is to be accessed. For some arrays, the array management service runs on a separate host; for other arrays it runs in the array itself.
- 4 Select **HP EVA** in the **Disk Array Host Type** pull-down menu.
- 5 Enter the user name and password for the array management service.
- 6 Clear (uncheck) the **Connect using port number** box.

Configuring a NetBackup policy for an HP EVA array method

Help is available for setting up a NetBackup policy in the NetBackup Administration Console.

See “Configuring a Snapshot Client policy” on page 52.

To configure a NetBackup policy for an HP EVA array method

- 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 an HP EVA method.
- 4** You can set the **Maximum snapshots (Instant Recovery only)** parameter. The maximum value is 16 snapshots or vsnaps per source device.

See “Maximum Snapshots parameter” on page 101.

HP EVA restrictions

Note the following array configuration restrictions. In essence, you cannot use two or more EVA snapshot methods for a given source disk.

Table 9-6

Array	Restrictions
VSnapshots	<p>Note the following restrictions:</p> <ul style="list-style-type: none"> ■ If any VSnapshots exist, you cannot create Snapshots or Clones of the source until existing VSnapshots are deleted. ■ A maximum of 16 VSnapshots can exist for any VDisk. ■ You cannot perform a rollback restore from a VSnap or a Snapshot. Rollback only works with clones.
Snapshots	<p>Note the following restrictions:</p> <ul style="list-style-type: none"> ■ If any Snapshots exist, you cannot create VSnapshots or Clones of the source until existing Snapshots are deleted. ■ A maximum of 16 snapshots can exist for any VDisk. ■ You cannot perform a rollback restore from a VSnap or a Snapshot. Rollback only works with clones.
SnapClones	<p>Any number of clones can be created for a VDisk, as long as disk space exists in the disk group.</p>

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 9-7 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 in the NetBackup Administration Console.
- 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 DSCSI 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 10000000C956A9B4 Sun          SUN - Solaris          0 V11
all
oigtaix03     0022 10000000C969F60E pSeries      IBM pSeries - AIX      0 V46
all
oigtaix02     0023 10000000C94AA677 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
```

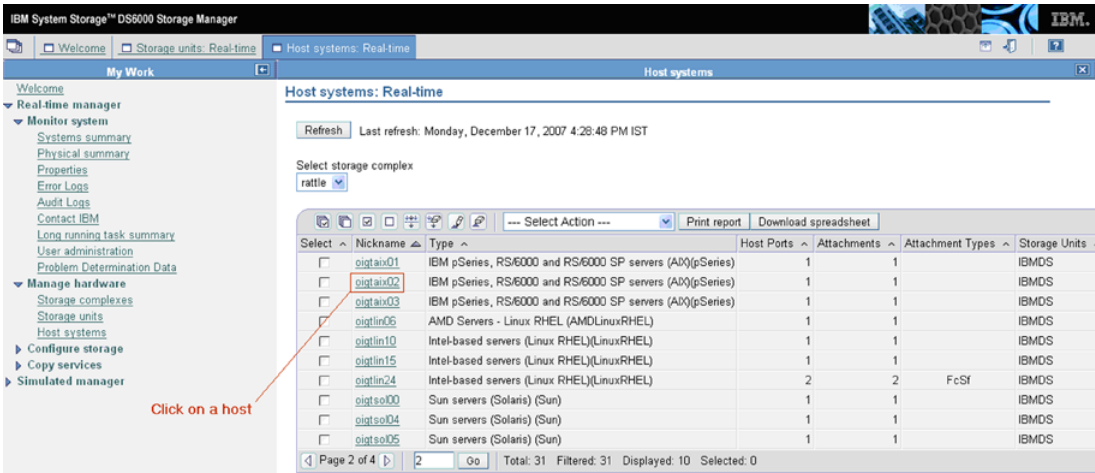
The values listed for `Vols` are the LUN ids.

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

```
/usr/openv/netbackup/bin/nbfirescan
```

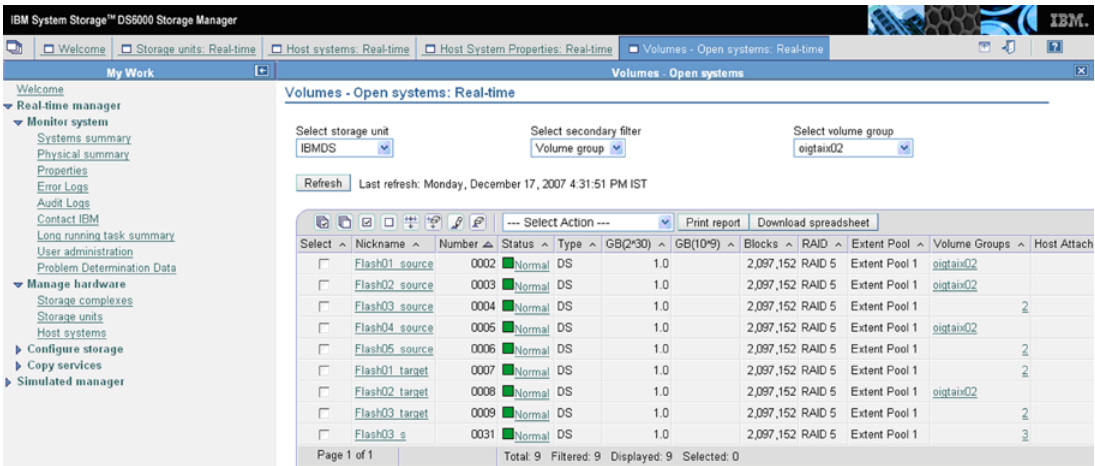
To use the IBM Storage Manager web interface to obtain the device identifiers

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



- 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.



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 192.

- 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 9-8 describes the snapshot method for the IBM DS4000.

Table 9-8 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 9-9 describes the tasks.

Table 9-9 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 197.
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 9-10 Software that is required for IBM 4000

Software	Where to install	Version
SMclient	Default location on NetBackup client: <code>/opt/IBM_DS4000/</code>	9.60 or higher
SMruntime	Default location on NetBackup client: <code>/opt/IBM_DS4000/</code>	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/openv/netbackup/bin/nbfirescan
```

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	<p>The maximum value is 4 snapshots per source device.</p>
Repository % of Base (100 for Instant Recovery)	<p>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.</p> <p>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).</p> <p>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%).</p> <p>For more details about the repository logical drive, refer to the IBM System Storage DS4000 Series and Storage Manager document.</p>

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 9-11 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: <code>/usr/lib/RMLIB/bin/whattrmver</code> See “Determining the current RAID Manager version” on page 201.

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/whattrmver
```

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 202.

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 202.

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/openv/netbackup/bin/nbfirescan
```

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 203.

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 9-12 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/whattrmver</code> See “Determining the current RAID Manager version” on page 201.

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 207.

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/openv/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/openv/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 207.

- 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:

- See “Important disk array method notes and restrictions” on page 148.
- 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`.

See “About Solaris `sd.conf` file” on page 155.

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 9-13 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. See “Configuring NetBackup to access the CLARiiON array” on page 166.
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 9-13 Issues with NetBackup and EMC CLARiiON arrays (*continued*)

Issue	Explanation/Recommended Action
Backups fail and the following message appears in the bpfis log: emccлариionfi: No more available HLU numbers in storage group. LUN <i>LUN number</i> cannot be LUN masked at this time	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: emccлариionfi: Could not find LUN <i>LUN number</i> in clonegroup <i>clonegroup name</i>	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: emccлариionfi: CLIDATA: Error: snapview command failed emccлариionfi: CLIDATA: This version of Core Software does not support Snapview	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: execNAVISECCLI: CLI Command [<i>CLI command</i>] failed with error [<i>error number</i>]	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 9-13 Issues with NetBackup and EMC CLARiiON arrays (continued)

Issue	Explanation/Recommended Action
<p>Policy validation fails for Standard policy with the following message:</p> <p>The error message is <code>Incorrect snapshot method configuration or snapshot method not compatible for protecting backup selection entries.</code></p>	<p>Policy validation for a Standard policy created with the <code>EMC_CLARiiON_Snapview_Snapshot</code> fails with error 4201.</p> <p>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.</p> <p>To fix the policy validation, add the following entry into the <code>emcclariionfi.conf</code> file:</p> <pre>[CLI_TOOL_INFO] "FILEPATH_NAVISEC_EXE"="<<NAVI CLI path>" "FILENAME"="<<NAVI CLI tool file name>"</pre> <p>For example, for the HP platform the path would be:</p> <pre>/usr/opensv/lib/vxfi/configfiles/emcclariionfi.conf</pre> <p>The entry would be:</p> <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 9-14 Issues with NetBackup and EMC Symmetrix arrays

Issue	Explanation/Recommended Action
<p>Point in time rollback fails and the following message appears in the <code>bpfis</code> log:</p> <p><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></p>	<p>See “Verifying that the clone is complete before doing a point in time rollback” on page 181.</p>

Table 9-14 Issues with NetBackup and EMC Symmetrix arrays (*continued*)

Issue	Explanation/Recommended Action
<p>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:</p> <p>An internal Snap or Clone error has occurred. Please see the symapi log file</p>	<p>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.</p>
<p>EMC_TimeFinder_Mirror backups fail and the following message appears in the bpfis log:</p> <p>emcsymfi: Invalid STD-BCV pair state</p>	<p>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.</p> <p>See "Fully synchronizing STD/BCV mirror pairs" on page 180.</p>
<p>TimeFinder backups fail and a message similar to the following appears in the bpfis log:</p> <p>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.</p> <p>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.</p>	<p>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.</p>

Troubleshooting NetBackup and HP EVA arrays

Table 9-15 Issues with NetBackup and HP EVA arrays

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

Table 9-15 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: client/server handshaking failed (26) " for HP_EVA_Snapclone FIM	Uninstall the VxVM software from the client.

See “Configuring NetBackup to access the EVA array” on page 188.

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 190.

See “Configuring array access for NetBackup hosts not named in a policy” on page 191.

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

Note these possible explanations:

Table 9-16 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.
<p>The snapshot device (clone) is also a source device in another device pair.</p> <p>The following message may appear in the <code>/usr/openv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log:</p> <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.
<p>The snapshot device (clone) and source device are not of equal size.</p> <p>The following message may appear in the <code>/usr/openv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log:</p> <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.
<p>The source device is already recording enabled for FlashCopy.</p> <p>The following message may appear in the <code>/usr/openv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log:</p> <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 9-16 Issues encountered with Snapshot (NetBackup status code 156)
(continued)

Issue	Explanation/Recommended Action
<p>The IBM FlashCopy license is not installed.</p> <p>The following message may appear in the /usr/openv/netbackup/logs/bpfis/ibmtsfi.log.date log:</p> <pre>CMUN03035E mkflash: Copy Services operation failure: feature not installed.</pre>	<p>Install the FlashCopy license on the storage subsystem. Contact IBM technical support or refer to your IBM array documentation.</p>
<p>The FlashCopy relationship is not recording enabled.</p> <p>The following message may appear in the /usr/openv/netbackup/logs/bpfis/ibmtsfi.log.date log:</p> <pre>CMUN03027E resyncflash: FlashCopy operation failure: action prohibited by current FlashCopy state. Contact IBM technical support for assistance</pre>	<ul style="list-style-type: none">■ 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.
<p>A FlashCopy relationship does not exist.</p> <p>The following message may appear in the /usr/openv/netbackup/logs/bpfis/ibmtsfi.log.date log:</p> <pre>CMUN03027E resyncflash: FlashCopy operation failure: action prohibited by current FlashCopy state. Contact IBM technical support for assistance.</pre> <p>A resynchronize operation was attempted on a FlashCopy pair that does not exist.</p>	<p>Verify that a FlashCopy pair does not exist, and then re-execute the backup.</p>

Table 9-16

Issues encountered with Snapshot (NetBackup status code 156)
(continued)

Issue	Explanation/Recommended Action
<p>Incremental Copy is in progress.</p> <p>The following message may appear in the <code>/usr/openv/netbackup/logs/bpfis/ibmtsfi.log.date</code> log:</p> <pre>CMUN02498E resyncflash: The storage unit is busy.</pre> <p>This message can appear for the following reasons:</p> <ul style="list-style-type: none">■ 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.	<p>Allow the background copy or maintenance activity to complete. Then re-run the backup.</p>

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 9-17 Explanations and recommended actions for status code 156

Problem	Description and recommended action for status code 156
<p>The array does not have enough free space.</p>	<p>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.</p> <p>The following messages may appear in the /usr/opensv/netbackup/logs/bpfis log:</p> <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 ibmdsf4700.</pre> <p>In addition, the following messages may appear in /usr/opensv/netbackup/logs/bpfis/ibmdsf.log.date:</p> <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> <p>Recommended action: Make sure that the array has enough space available for the snapshot.</p>
<p>The maximum number of FlashCopies (4) already exist on the array.</p>	<p>The following messages may appear in /usr/opensv/netbackup/logs/bpfis/ibmdsf.log.date:</p> <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> <p>Recommended action: Delete any FlashCopies that NetBackup did not create.</p>
<p>The Access Logical Drive is not mapped for the NetBackup client or alternate client at LUN 31 or 7.</p>	<p>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.</p> <p>Recommended action: Create and map an Access Logical Drive. Contact IBM technical support or refer to your IBM array documentation.</p>

Table 9-17 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 201.
- 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 204.

NetBackup policy validation fails with Hitachi arrays

Note these possible explanations:

Table 9-18 NetBackup policy validation fails

Issue	Explanation/Recommended Action
<p>Look for the following error in the <code>/usr/opensv/netbackup/logs/bpfis/hitachi.log.<date></code> log:</p> <pre>Library RMLIB init failed</pre>	<p>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.</p>
<p>The Hitachi command device is not unmasked. See the sample log messages in the next row.</p>	<p>Refer to Hitachi documentation for creating and unmasking command devices.</p>
<p>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.</p> <p>The <code>/usr/opensv/netbackup/logs/bpfis/hitachi.log.<date></code> log may contain messages similar to the following:</p> <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/rdisk/c4t50060E801029F700d2s6" 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> <p>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).</p>	<p>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.</p> <p>To determine if the command device is recognized by the operating system, try device discovery commands such as the following:</p> <pre>devfsadm cfgadm -al</pre> <p>A log message for the enclosure ID would include an entry such as the following:</p> <pre>c3t50060E801029F700d28 <HITACHI-DF600F-CM-0000 cyl 52 alt 2 hd 50 sec 768></pre> <p>which shows that the device is visible as <code>c3t50060E801029F700d28</code>.</p>
<p>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.</p> <p>See the sample log messages in the next bullet.</p>	<p>Specify the correct snapshot method or snapshot devices.</p>

Table 9-18 NetBackup policy validation fails (*continued*)

Issue	Explanation/Recommended Action
<p>A disk pair was not created for the source device and snapshot device specified in the NetBackup policy's Snapshot Resources.</p> <p>The <code>/usr/openv/netbackup/logs/bpfis/hitachi.log.<date></code> log may contain messages similar to the following.</p> <p>If the snapshot method is Hitachi_CopyOnWrite:</p> <pre>Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> HITACHI_FIM_SNAPSHOT not supported for 10</pre> <p>If the snapshot method is Hitachi_ShadowImage:</p> <pre>Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> HITACHI_FIM_SHADOW_IMAGE not supported for 10</pre>	<p>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.</p>
<p>In the policy's Snapshot Resources, the device identifier for the source device or snapshot device is invalid.</p> <p>The <code>/usr/openv/netbackup/logs/bpfis/hitachi.log.<date></code> log may contain messages similar to the following:</p> <pre>Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> gettrminfo failed. Fri Mar 21 2008 16:26:49.173893 <Pid - 9477 / Thread id - 1> operation failed with error number <> with message <msg>'.</pre> <p>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:</p> <pre>Mon May 12 2008 21:32:32.088876 <Pid - 8040 / Thread id - 1> gettrminfo is called for '9999'. Mon May 12 2008 21:32:32.089736 <Pid - 8040 / Thread id - 1> gettrminfo 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>	<p>Recommended action: Make sure that the identifiers are correctly entered in the policy's Snapshot Resources. Specify source and snapshot IDs without leading zeros.</p> <p>See "Configuring a NetBackup policy for Hitachi_ShadowImage or Hitachi_CopyOnWrite" on page 204.</p>

Table 9-18 NetBackup policy validation fails (continued)

Issue	Explanation/Recommended Action
The RAID Manager library <code>libsvrrm.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>libsvrrm.so</code> is not supported.	Recommended action: Look for the <code>Library RMLIB version</code> message in the <code>/usr/opensv/netbackup/logs/bpfiis/hitachi.log.<date></code> log. See “Determining the current RAID Manager version” on page 201.

Backup or restore or image expiration fails with Hitachi arrays

The following table lists possible explanations:

Table 9-19 Backup or restore or image expiration fails

Issue	Explanation/Recommended Action
	<p>These messages indicate that the instance number is not available. A problem may exist with the instance number management logic.</p> <ul style="list-style-type: none">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 9-19 Backup or restore or image expiration fails (continued)

Issue	Explanation/Recommended Action
<p>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.</p> <p>The</p> <pre>/usr/opensv/netbackup/logs/bpfis/hitachi.log.<date></pre> <p>log contains the following message:</p> <pre>Couldn't get instance no failed with message</pre> <p>The log may contain the following additional messages:</p> <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> <p>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.</p> <p>Note also the following potential messages:</p> <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 9-19 Backup or restore or image expiration fails (*continued*)

Issue	Explanation/Recommended Action
<p>16:26:49.818308 <Pid - 9477 / Thread id - 1> Exiting Function attachCmd [156, providers/hitachi/hitachi_rmlibintf.cpp]</p> <p>Another example message is the following:</p> <p>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'</p>	
<p>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.</p>	<p>Recommended action: Make sure that the clone (or snapshot) device has the same default controller as the source device. See the Hitachi documentation.</p>

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
- About restores from a FlashBackup backup
- Restoring a large number of files in a clustered file system (VxFS on UNIX Only)
- 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	<p>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.</p> <p>See the NetBackup Administrator's Guide, Volume I, for instructions on doing manual backups.</p>
User-directed backup and archive	<p>From a NetBackup client, the user can execute a Snapshot Client backup. The NetBackup administrator must configure an appropriate snapshot policy with schedule.</p> <p>See the NetBackup Backup, Archive, and Restore Getting Started Guide for instructions on doing user-directed backups and archives.</p>

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.

About restores from a FlashBackup backup

Using the Backup, Archive, and Restore interface, you can restore individual directories or files (or an entire raw partition) from a FlashBackup 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.

Note: In the Backup, Archive, and Restore interface, set the policy type to **FlashBackup** for UNIX clients and **FlashBackup-Windows** for Windows clients.

Note the following before you start the restore:

- A FlashBackup (or FlashBackup-Windows) backup supports both individual file restore and raw partition restore. You can do either type of restore from the same backup. To restore individual files, select **Normal Backups** on the **Restore Files** tab; to restore an entire raw partition, select **Raw Partition Backups**.
- To restore a raw partition, you must have administrator capabilities on the NetBackup server.

- An entire raw partition can be restored from a full backup only. FlashBackup incremental backups only support individual file restores.
- Ensure that the device file for the raw partition exists before the restore.
- The overwrite option must be selected for raw partition restores. The device file must exist and the disk partition is overwritten.
- To restore a very large number of files (when individual file restore would take too long), you can do a raw-partition restore. Redirect the restore to another raw partition of the same size and then copy individual files to the original file system.

Notes for FlashBackup and UNIX client restore

Note the following for UNIX clients:

- To restore an entire raw partition, ensure that the partition is not mounted and not in use. (For this reason, you cannot perform a raw partition restore to the root partition, or to the partition on which NetBackup is installed.) If a database uses the partition, shut down the database. The partition must be the same size as when it was backed up; otherwise, the results of the restore are unpredictable.
- After a raw partition restore of a VxFS file system, a file system consistency check (fsck) is usually required before the file system can be mounted.

Notes for FlashBackup and Windows client restore

Note the following for Windows clients:

- For raw partition restores, you must select the overwrite option. The device file must exist and the disk partition is overwritten during the restore.
- For raw partition restores, ensure that the partition is mounted (designated as a drive letter) but not in use. (For this reason, you cannot perform a raw partition restore to the root partition, or to the partition on which NetBackup is installed.) If a database uses the partition, shut down the database. The partition must be the same size as when it was backed up; otherwise, the results of the restore are unpredictable.
- For raw partition restores, make sure to select the drive letter that is in the same format as was specified in the FlashBackup-Windows policy (for example, \\.\E:). If you select E:\, the restore fails.

Restoring a large number of files in a clustered file system (VxFS on UNIX Only)

In a clustered file system, to restore a large number of files (such as 100,000 or more), the restore finishes sooner if the file system is mounted on a local host. The file system can then be mounted as a shared mount after the restore.

To restore a large number of files in a clustered file system

- 1 Stop all applications (on any nodes) that use the file system.
- 2 Unmount the file system.
- 3 Mount the file system locally.
- 4 Perform the restore.
- 5 Share the mounted file system again, and restart applications (if any).

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 several restore features unique to Instant Recovery that require special instructions.

Block-level restore (for VxFS_Checkpoint snapshots)	See "About Instant Recovery: block-level restore" on page 234.
File promotion (for VxFS_Checkpoint or NAS_Snapshot snapshots)	See "About Instant Recovery: file promotion" on page 235.
Fast File Resync for Windows (for VxVM and FlashSnap snapshots)	See "About Instant Recovery: Fast File Resync (Windows clients only)" on page 236.
Rollback (for VxFS_Checkpoint, VxVM, VSS, FlashSnap, or NAS_Snapshot snapshots, OST_FIM, and the disk array methods)	See "About Instant Recovery: point in time rollback" on page 238.

About Instant Recovery: block-level restore

If the Instant Recovery snapshot was made with the VxFS_Checkpoint method, large files can be recovered faster by means of block-level restore. Only the blocks that have changed are moved from the snapshot to the client's primary fileset.

Note the following:

- Block-level restore requires the VxFS File System.
- Block-level restore is available only when restoring files to the original location on the client, AND when the snapshot method for the backup was VxFS_Checkpoint.
- If the snapshot method for the backup was VxFS_Checkpoint and the files to be restored are in an Oracle database, block-level restore is automatically enabled.

More information is available about activating and deactivating block-level restore. See “Activating and deactivating block-level restore” on page 235.

Activating and deactivating block-level restore

To activate block-level restore

- ◆ Create the following (empty) file on the client:

```
/usr/opensv/netbackup/PFI_BLI_RESTORE
```

After this file is created, all subsequent restores of the client's data use block-level restore.

To deactivate block-level restore

- ◆ Delete (or rename) the `PFI_BLI_RESTORE` file.

When block-level restore is activated, it is used for all files in the restore. Block-level restore may not be appropriate for all of the files. It may take longer to restore a large number of small files, because they must first be mapped.

About Instant Recovery: file promotion

If the Instant Recovery snapshot was made on UNIX with either VxFS_Checkpoint or NAS_Snapshot, large files that have had many changes since the backup can be recovered more quickly by means of file promotion. File promotion optimizes single-file restore by using a minimum of I/O to recover the files.

Notes on file promotion

Only regular files can be promoted, not file links or directories.

Note the following regarding VxFS_Checkpoint:

- File promotion requires the VxFS File System version 4.0 or later.
- File promotion can be done only from the last Instant Recovery snapshot that was made with the VxFS_Checkpoint method.

- File promotion is available only when restoring files to the original location on the original client.

Note the following regarding NAS_Snapshot:

- File promotion is available when restoring to the original volume on the original client.
- File promotion can be done from older snapshots, but any newer NAS snapshots are deleted after the file promotion takes place.
- The file system requirements depend on the NAS vendor.
- For further requirements specific to your NAS vendor, see the *NetBackup for NDMP Supported OS and NAS Appliance Information* online document. That document can be accessed from the following:
<http://www.veritas.com/docs/000027113>

About file promotion

The file promotion procedure is the same as the standard restore procedure for NetBackup.

See the NetBackup Backup, Archive, and Restore Getting Started Guide.

No special settings or choices are required when you are using the Backup, Archive, and Restore interface.

If the requirements are met, NetBackup automatically attempts file promotion for the file.

See “Notes on file promotion” on page 235.

Otherwise the restore of the file takes place in the standard manner, without file promotion. All file data is copied from the snapshot to the primary file system. The NetBackup progress log indicates how many files were promoted and how many files that could not be promoted were restored by means of tar.

About Instant Recovery: Fast File Resync (Windows clients only)

If the Instant Recovery snapshot was made on Windows with VxVM or FlashSnap, large files that have had many changes since the backup can be recovered more quickly by means of Fast File Resync. (Fast File Resync is a type of file promotion.) Only the blocks that have changed are moved from the snapshot to the client's primary fileset.

Notes on Fast File Resync (FFR)

Note the following:

- FFR requires Storage Foundations for Windows 4.1 or later and the licensed FlashSnap option.
- FFR can be done only from an Instant Recovery snapshot that was made with the VxVM or FlashSnap method.
- FFR is available only when you restore to the original location on the original client.
- The **overwrite existing files** option must be selected.

Notes on Fast File Resync (FFR) and the files to be restored

Note the following:

- The original files must be present on the client and are overwritten by the restore.
- The names and creation times of the original files and the snapshot files must be identical.
- Files must be larger than 20 MB and be formatted in NTFS.
- Files must not be compressed or encrypted.
- There must be no open handles on either the original file or the snapshot file.

About using Fast File Resync

The Fast File Resync procedure for restoring individual files is the same as the standard restore procedure for NetBackup. Refer to the NetBackup Backup, Archive, and Restore Getting Started Guide. No special settings or choices are required when you use the Backup, Archive, and Restore interface.

When NetBackup uses Fast File Resync

If the requirements are met, NetBackup automatically attempts Fast File Resync first.

See “Notes on Fast File Resync (FFR)” on page 236.

If Fast File Resync cannot be used for a file, the restore of that file takes place in the standard manner, without Fast File Resync. That is, all file data is copied from the snapshot to the primary file system. The NetBackup progress log indicates how many files were re-synchronized, and how many files that could not be re-synchronized were restored by means of tar.

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.

- Rollback can be done only from the backups that were enabled for Instant Recovery and made with one of the following methods: VxFS_Checkpoint, VxVM, FlashSnap, NAS_Snapshot, or the disk array methods.
- If the backup was made with the EMC_TimeFinder_Clone method and the clone is not fully created, a rollback cannot succeed.
To verify that the clone is complete before you do a rollback:
See “Verifying that the clone is complete before doing a point in time rollback” on page 181.
- For the backups that were made with the VxFS_Checkpoint method, rollback requires the VxFS File System 4.0 or later and Disk Layout 6. For NAS_Snapshot, the file system requirements depend on the NAS vendor.
- Rollback deletes any VxFS_Checkpoint snapshots or NAS_Snapshot snapshots (and their catalog information) that were created after the creation-date of the snapshot that you restore.
- If the primary file system is mounted and the snapshot resides on a disk array, the rollback attempts to unmount the file system. Any I/O on the primary device is forcibly stopped if the unmount succeeds. To be safe, make sure that no I/O occurs on the primary device before a rollback.
If the attempt to unmount the primary file system fails, the rollback does not succeed. You should halt I/O on the device and retry the rollback. For example, if a terminal session has accessed the file system through the `cd` command, change to a directory outside the file system and retry the rollback.
- Rollback is available only when you restore the file system or volume to the original location on the client.

- When a file system rollback starts, NetBackup verifies that the primary file system has no files that were created after the snapshot was made. Otherwise, the rollback aborts.
- Rollback from OST_FIM type snapshot can be done from copy one only.
- For the rollback from OST_FIM type snapshot, refer to the NetBackup Replication Director Solutions Guide.

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/jbpSA &
```

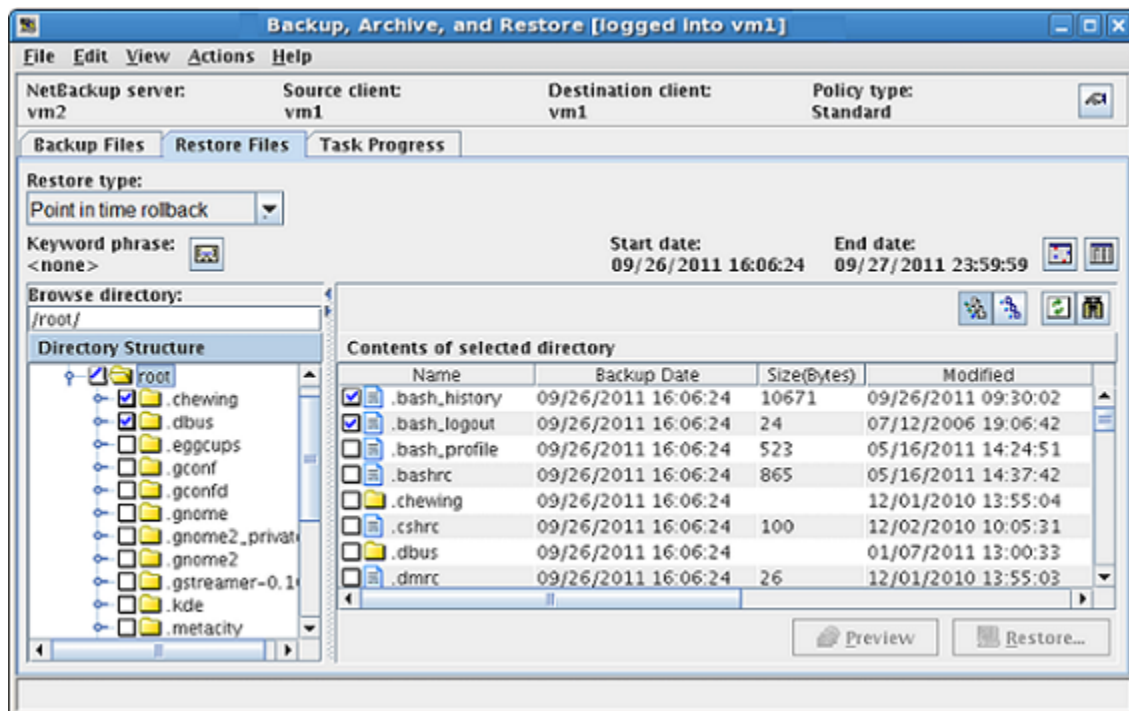
- 2 Click the **Restore Files** tab.
- 3 Click **Actions > Specify NetBackup Machines** to specify the server, source client, policy type, and destination client.
- 4 For the Restore Type, select **Point in Time Rollback**.

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.

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 238.

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.

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.

To perform snapshot rollback (Windows)

- 1 Start the Backup, Archive, and Restore interface.


Click **Start > Programs > Veritas NetBackup > Backup, Archive, and Restore**.

- 2 From the **Select for Restore** drop-down list, select **Restore from Point in Time Rollback**.

- 3 Click **File > Specify NetBackup Machines and Policy Type** to specify the server, source client, policy type, and destination client.

- 4 In the **NetBackup History** pane, click the backup image to restore.

Only Instant Recovery backups are displayed in the **NetBackup History** pane, for all dates (you cannot set a date range).



You can select root level or mount points (file systems or volumes), but not folders or files at a lower level.

- 5 In the **All Folders** pane, 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.
- 6 Click **Actions > Start Restore of Marked Files**.
The only destination option is **Restore everything to its original location**.
- 7 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 238.

If the checks do not pass, the rollback aborts and the progress log states that rollback could not be performed because file verification failed.

The remainder of the procedure is identical to a normal restore as explained in the NetBackup Backup, Archive, and Restore Getting Started Guide.

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 <code>FORCE_RESTORE_MEDIA_SERVER</code> option in <code>bp.conf</code> 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 246.

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 **Keep 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 246.

- 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,MNTPOINT=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
```


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 Snapshot Client
- Snapshot Client 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

- 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
- An NDMP snapshot policy for wildcard backup fails with error 4201
- Troubleshooting with bpfis log
- Limitations of using HP-UX 11.31

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 12-1 on page 255., 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, Snapshot Client 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 12-1 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, Snapshot Client messages are logged to the following directories on the master server. Create these directories using an access mode of 755.

Table 12-2 UNIX logging directories for restore

Path of log directory	Where to create the directory
<code>/usr/opensv/netbackup/logs/bprestore</code>	NetBackup master server
<code>/usr/opensv/netbackup/logs/bprd</code>	NetBackup master server

Table 12-2 UNIX logging directories for restore (continued)

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bpbrm	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, Snapshot Client messages are written to the folders that are listed in See Table 12-3 on page 257. 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, Snapshot Client 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 12-3 Windows logging folders for backup

Path of log directory	Where folder is created
<i>install_path\NetBackup\logs\bpbm</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\bpffi</i>	NetBackup client or alternate client
<i>install_path\NetBackup\logs\bpbk</i>	NetBackup client or alternate client
<i>install_path\NetBackup\logs\bpfi</i>	NetBackup client or alternate client

Windows logging folders for restore

During a restore, Snapshot Client messages are logged to the following folders on the master server.

Table 12-4 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\bpri</i>	NetBackup master server
<i>install_path\NetBackup\logs\bpbm</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. For details on the Logging Assistant, see the NetBackup Administrator's Guide, Volume I.

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 12-5 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 12-5 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 12-5 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 12-5 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.

Table 12-5 VxMS logging levels (continued)

Level	Description
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 Snapshot Client 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 Snapshot Client

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 52.
- 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-tl4-0`) 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/opensv/netbackup/db/config/SIZE_DATA_BUFFERS_TPC.policy_name
/usr/opensv/netbackup/db/config/SIZE_DATA_BUFFERS_TPC.storage_unit_name
/usr/opensv/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.

Snapshot Client 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 Snapshot Client 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:

```
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 90.

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/openv/netbackup/bin/driver/snaplist
```

- 2 For each snapshot that is listed, run the following to make sure a `bpbkar` 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 `bpbkar` process by entering the following:

```
/usr/openv/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 Snapshot Client 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/openv/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 281.

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 `fsckptadm list` command, is the following:

```
/vm2
```

Example entry:

```
/usr/lib/fs/vxfs/fsckptadm 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/fsckptadm 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 270.
- 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:

```
vxvg list
```

- The format of the disk group name is as follows:

```
SPLIT-primaryhost_diskgroup
```

- Deport the VxVM disk group:

```
vxvg -g split_diskgroup deport
```

- On the primary (original) client, import and join the VxVM disk group:

```
vxassist rescan
vxvg -g split_diskgroup import
vxvg -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/opensv/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/opensv/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.

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 **PFtest123456789vol** and the suffix **00043c8aaa** is added to it, the volume name exceeds the limit. The command `vxassist snapshot` does not recognize the name **PFtest123456789vol_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:

Snapshot creation fails when the same volume is mounted on multiple mount points of the same host**To verify if the zoning is correct**

- 1 Check the FC/scsi adapter on the host in `/usr/sbin/lssdev -C -c adapter -Sa -F 'name'`

The Output should be:

fcs0

fcs1

- 2 Check FC SCSI I/O Controller in `/usr/lib/methods/cfgefc -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/cfgefscsi -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`

`rsz=32768,wsz=32768,NFSv3,dev=4000033` and `/test1` on

`f3170-7-15:/vol/sample1 rsz=32768,wsz=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 Snapshot Client 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.

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 255.

See “Windows logging folders for backup” on page 256.

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 7.5 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
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.

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 127.

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 125.

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 Snapshot Client 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 285.

See “Identifying and removing a left-over snapshot” on page 264.

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/openv/netbackup/bin/driver/snapon /var /dev/rdisk/c2t0d0s3
```

Example 2:

```
/usr/openv/netbackup/bin/driver/snapon /dev/vx/rdsk/omo/tcpl  
/dev/vx/rdsk/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/openv/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/rdsk/omaha/tcpl					
	cache = /dev/vx/rdsk/omaha/sncache					
16	7857	2097152	128	0	0	05/29/03 12:16:00
	device = /dev/vx/rdsk/omaha/vol01					
	cache = /dev/vx/rdsk/omaha/sncache					
17	7908	20971520	4224	0	0	05/29/03 12:17:38
	device = /dev/vx/rdsk/omaha/vol03					
	cache = /dev/vx/rdsk/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	<p>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.</p> <p>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.</p>
cached	<p>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.</p>
minblk	<p>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>.</p>
err	<p>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>snapoff</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>.</p>
time	<p>The time at which the snapshot was started.</p>
device	<p>The raw partition containing the client's file system data to back up (snapshot source).</p>
cache	<p>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.</p> <p>To determine the total space that is used in a particular cache partition by all active snapshots, use the <code>snapcachelist</code> command.</p>

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	<p>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.</p> <p>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.</p>

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 Snapshot Client.

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.

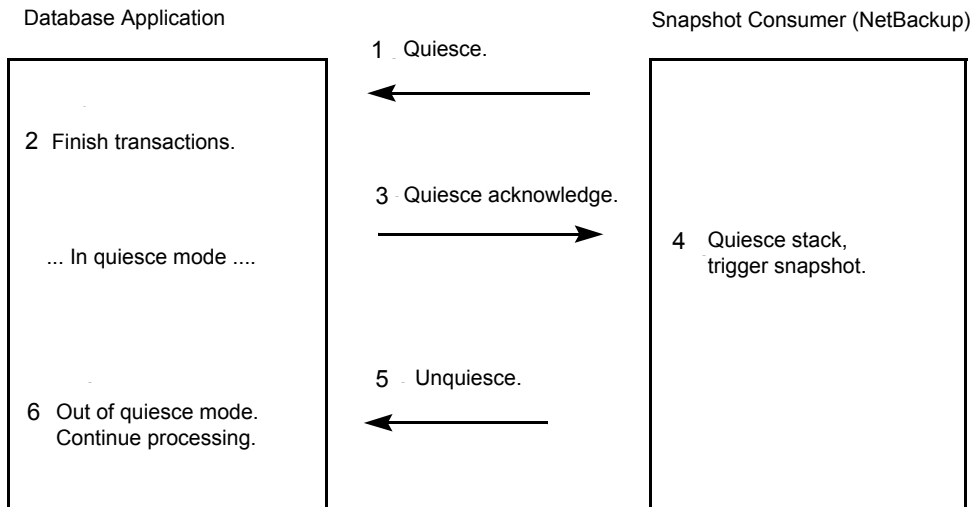
About quiescing the database application

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

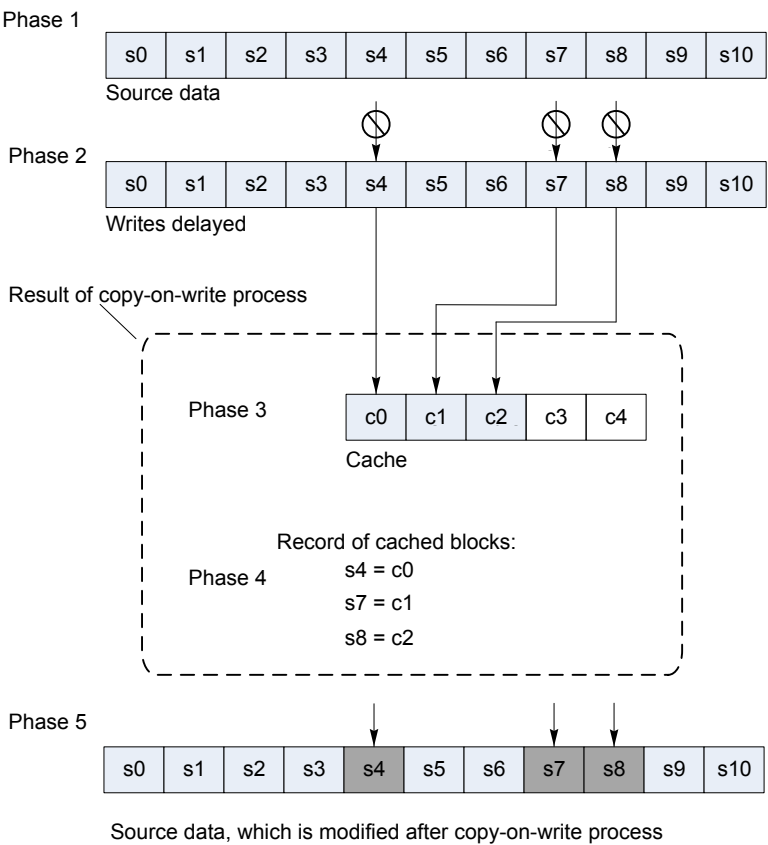
A copy-on-write snapshot is a detailed account of data as it existed at a certain moment. Unlike a mirror, a copy-on-write is not a copy of the data, but a specialized account of it.

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).

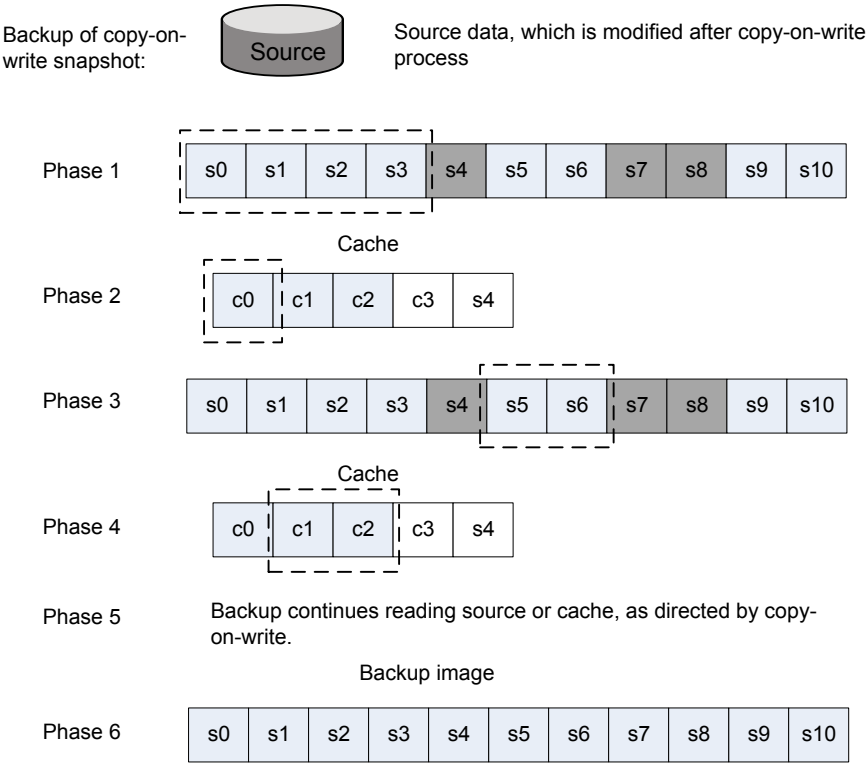
Phase	Action
Phase 3	Copy-on-write process writes contents of blocks s4, s7, and s8 to cache. These blocks write to cache only once, no matter how many times they change in the source during the snapshot.
Phase 4	Copy-on-write process keeps a record of the number of writes to cache.
Phase 5	Write requests 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 260
3pc.conf file 43

A

abnormal termination 264, 282
access time not updated 78
activate block-level restore 235
Active Directory 82, 140
 and FlashBackup 82
actual device file name 86–87
Administrator account and NAS_Snapshot 119
AIX
 and VxFS 45
 media servers, restrictions 38
ALL_LOCAL_DRIVES entry 36, 55, 104, 228
Allow multiple data streams 77
alternate client
 defined 38
Alternate client backup
 configuring 73
alternate client backup 20, 26, 58, 85
 and FlashSnap 133
 and split mirror 28
 and VVR 136
 introduction 27
 requirements 73
 restrictions 73
 testing setup 133, 137, 140
Any_available storage unit 59
APP_NBU_VVR 137
arbitrated loop 39
archive bit
 incremental backup 73
archives 232
auto option (Provider Type for VSS) 68
auto snapshot selection 59, 62, 74, 102
automatic backup 231

B

Backup
 Archive
 Restore 239
backup
 agent 19, 38
 automatic 231
 local 24, 33
 logs 255–257
 manual 232
 off-host
 configuration 55, 57
 prerequisites 227
 raw partition 56, 228
 retention period
 NAS_Snapshot 138
 scripts 71
 techniques (overview) 21
 types supported 86
 user-directed 232
backup agent 57
backup retention level 101
Backup Selections list 55
 ALL_LOCAL_DRIVES entry 36, 55, 104, 228
 and Instant Recovery 104
 block vs. character device 228
 directives 90
 FlashBackup 87
 symbolic link 77
BLIB 52
block device file (vs character) 228
block level incremental backup 52
block-level restore 234
 how to activate 235
 restriction 235
bp.conf file 254
bpbkar
 log 255
 process 264
bpbrm log 255–256
bpdgclone command 271

- bpfis
 - command 247, 250, 265
 - log 255
- bppfi log on client 255
- bprd log 255
- bprestore log 255
- bptm log 255–256
- BusinessCopy method 262

C

- cache 65
 - definition of 39
 - flushing for quiesce 290
 - object 108, 131
 - overflow 282
 - partition 90
 - partition out of space 263
 - requirements 124
 - size 282
 - specifying raw partition for 124
- CACHE= directive 88, 90, 127
- Change Policy dialog 61, 72
- character device file (vs block) 124, 228
- checking logs 254
- checkpoint
 - removing (VxFS) 266
 - VxFS file system 128
- CLARiiON
 - fractured (split) clone after rollback 211
- client data
 - prerequisites for off-host backup 227
- client list 87
- client read timeout 229
- clone
 - fractured (split) after rollback 211
 - removing 266, 270
 - VxVM disk group 270
- clone group (CLARiiON) 169
- clone private LUNs (CLARiiON) 169
- Cluster Volume Manager. *See* CVM
- CommandCentral 59
- compression 52
- configuration
 - auto method 59
 - backup selections 87
 - client list 87
 - prerequisites 51
 - procedure 53, 60, 102
 - snapshot parameters 65

- configuration *(continued)*
 - supported data types 82
- Configure snapshot policy 52
- copy manager (see third-party copy) 42
- copy-on-write
 - defined 39
 - overview 22, 290
- cp command 246
- credentials
 - disk array
 - rollback and VSS 158
- cross mount points (disabled) 77
- customer support 260
- CVM
 - about support for 141
 - enabling vxvm or FlashSnap snapshots in 142
 - execute VxVM commands remotely 143

D

- data change object 107
- data consistency (quiesce) 288
- data mover 43
 - Network Attached Storage 58
- data mover (see third-party copy) 42
- Database Edition for Oracle 35
- database files (Windows) 83, 140
- DB2 94
- DCO plexes 134
- deleting
 - a clone 270
 - snapshots 100
- deport
 - disk group 133, 135
 - snapshot 248, 251
- device
 - serialization 38, 228, 261
- Device Configuration wizard 52
- device file name (FlashBackup) 87
- differential option (Snapshot Attribute for VSS) 69
- directives
 - for scripts 71
- directives (Backup Selections list) 90
- directories for troubleshooting 254
- disabling snapshots 79
- disk
 - restoring snapshot from 246
 - SCSI vs. IDE 261
 - visibility 261

- disk array methods
 - EMC CLARiiON arrays 162
 - EMC disk groups for VSS differential snapshots 159
 - Hitachi 201
 - HP EVA 184
 - HP-XP 206
 - IBM DS4000 196
 - IBM DS6000 and DS8000 190
 - troubleshooting 209
 - verify NetBackup access to arrays 156
 - verify that VSS can make a snapshot 160
- disk array snapshot methods
 - about 145
 - advantages 147
 - dynamic multipathing 154
 - HBAs 154
 - Instant Recovery 148
 - list of 149
 - notes on 148
 - OS-specific configuration 154
 - persistent target bindings 154
 - Solaris sd.conf 155
 - tasks for admin 151
 - types 147
 - use in clusters 148
- disk arrays
 - and point in time rollback 238
- disk group
 - clone 270
 - shared in VxVM 124, 133
- disk layout for Storage Checkpoint 94
- disk restore
 - caution 249
- disk snapshot
 - restoring from 245–246
- disk storage units 59, 228
- DLL patch for VxFS 36
- DSCLI 190
- duplex mode and performance 261

E

- EMC CLARiiON
 - configure NetBackup policy for 175
 - device identifier 172
- EMC CLARiiON arrays 162
- EMC CLARiiON clone group 169
- EMC CLARiiON software requirements 163
- EMC CLARiiON storage group 167

- EMC Symmetrix 176
 - configuring NetBackup clients 178
 - software requirements 176
 - Solutions Enabler 176, 178
- EMC TimeFinder Snap 159
 - test 162
- EMC_CLARiiON_SnapView_Clone 168
- EMC_CLARiiON_SnapView_Snapshot 173
 - reserved LUN pool 174
- EMC_TimeFinder_Clone 181
 - and rollback 181, 238
- EMC_TimeFinder_Mirror 180
- EMC_TimeFinder_Snap 182
- encryption 52
- exceptions to exclude list 83
- exclude list 83, 87
- Extended Copy command 20, 26, 42, 58
- extent
 - defined 39

F

- Fast File Resync (in Storage Foundation for Windows 4.1) 102–103, 234, 236
- fast mirror resynch 132
- FastResync 39, 41, 102–103, 134
- FFR (Fast File Resync) 102–103, 234, 236
- Fibre Channel
 - defined 39
 - types supported 39
- file pathname (max length) 55, 76
- file promotion 20, 234–236
- file system
 - multi-volume 36, 128, 131
- file systems
 - defined 40
 - quiescing 290
- files list (see Backup Selections list) 55
- FlashBackup 281
 - actual device name 87
 - and VxFS_Checkpoint 128
 - backup selections 87
 - device file name 86
 - features 81
 - files to back up 86–87
 - restoring 232
- FlashSnap method 36, 62, 102, 237
 - and status 156 260
 - deporting the disk 133
 - preparing for alternate client backup 133

- FlashSnap method (*continued*)
 - restoring from disk image 248, 251
 - with VxVM shared disk group 133
- flush file system 290
- FORCE_RESTORE_MEDIA_SERVER option 244
- fsck (after raw partition restore of vxfs system) 233
- fsckptadm command 41, 266
- full-sized instant snapshot 65
- full-sized instant snapshots 109, 132–133

G

- gatekeeper devices 178–179
- get_license_key 47
- GID 73
- glossary of terms 38
- group of storage units 59

H

- hardware option (Provider Type for VSS) 69
- hardware-level disk restore 249
- Hitachi disk arrays 201
 - configure NetBackup for 202
 - configuring a NetBackup policy for 204
 - nbfirescan 202
 - obtain identifiers 203
 - pair status 202
 - RAID Manager version 201
 - software requirements 201
- Hitachi_CopyOnWrite
 - configuring a NetBackup policy for 204
- Hitachi_ShadowImage
 - configuring a NetBackup policy for 204
- HOMRCF 150
- HP 90
- HP EVA
 - configure NetBackup policy for 189
 - software requirements 185
 - VDisk 189
 - VSnapshots 189
- HP EVA disk arrays 184
- HP-XP disk arrays 206
 - configure NetBackup 207
 - nbfirescan 207
 - obtain identifiers 207
 - pair status 206
 - software requirements 206
- HP_XP disk arrays
 - configuring a NetBackup policy 208

- HP_XP_BusinessCopy
 - configuring a NetBackup policy 208
- HP_XP_Snapshot
 - configuring a NetBackup policy 208

I

- I/O components
 - supported 82
- IBC messages 137
- IBC send/receive timeout (seconds) 66
- IBM DS4000 disk array 196
 - configure a NetBackup policy for 200
 - configure for NetBackup 199
 - nbfirescan 198
 - SMclient 197
 - software requirements 197
 - verify NetBackup client access to 198
- IBM DS6000 and DS8000 disk arrays 190
 - configure array for NetBackup 191
 - configure NetBackup policy for 195
 - IBM_DiskStorage_FlashCopy 195
 - software requirements 190
- IBM_DiskStorage_FlashCopy 195
- IBM_StorageManager_FlashCopy 196
 - configure a NetBackup policy for 200
- IDE vs. SCSI 261
- identification descriptor 59
- importing disk group 135–136
- include list 87
- incremental backup
 - archive bit and timestamp 73
- Inline Tape Copies (Vault) 38
- inquiry page code 83 38, 228, 261
- installation
 - of Snapshot Client 47
- Instant Recovery 21
 - and HP EVA snapshot methods 150
 - defined 40
 - deleting snapshots 100
 - disk array snapshot methods 148
 - Fast File Resync 102–103, 234, 236
 - Schedule tab 103
 - selecting in policy 102
 - snapshot only 102
 - Snapshot Resources pane 100
 - volume name restriction 95, 106, 131
- instant snapshots 65, 108–109, 131–132

J

jbpSA 239

K

Keep snapshot after backup 66, 264
 restoring from image 245–246
 kernel messages 256

L

left over snapshot 264, 282
 license keys
 installing 47
 Limit jobs per policy 77
 limitations 36
 links (in Backup Selections list) 77
 Linux
 and VxFS 45
 Local Host
 network configuration for 33
 local host backup method
 network configuration for 24
 local system account 119
 lock file system 290
 logging
 directories to create 254
 VxMS 257
 logical volume (as raw partition) 124
 logs 254–257
 creating for UNIX 255
 creating for Windows 256
 loop (Fibre Channel) 39
 ls command 41
 LVM 82

M

manual backup 232
 mapping
 defined 40
 Maximum jobs per client 77
 Maximum multiplexing per drive 77
 maximum pathname length 55, 76
 Maximum Snapshots (Instant Recovery only) 67, 101, 110, 138
 Media multiplexing 77
 media server (see NetBackup Media Server) 57–58
 messages file 256
 method
 selecting off-host backup 55, 57

method (*continued*)

 selecting snapshot 61

mirror 22

 access time 78
 compared to copy-on-write 24
 defined 40
 fast resynch 132
 overview 22
 preparing 106
 rotation 99
 VxVM snapshot 41, 131

mklogdir script 255

mklogdir.bat 256

modprobe.conf file (Linux) 156

mover.conf file 43

 and AIX 38

multi-volume system (VxFS) 36, 128, 131

Multiple Copies (Vault) 38

multiple data streams 77

 configuring 77

multiplexing 38, 228

MVS 36, 128, 131

N

NAS

 off-host backup 26

NAS filer

 as backup agent 58

NAS_Snapshot 21, 103–104, 138, 234

 access for NetBackup 119

 backup retention period 138

 licensing 118

 logging info 255–256

 name 122

 notes

 requirements 118

NAS_Snapshot method 62

naviseccli 165, 172

Navisphere 163

nbfirescan 156

NBU_CACHE 108

nbu_snap method 63, 123, 281

 with VxVM shared disk group 124

NDMP 21

 access web info 43

 licensing 118

NDMP host

 as backup agent 58

NDMP protocol version 119

- NDMP snapshot 138
- ndmp unified log (VxUL) 255–256
- NDMP V4 103
- NetBackup Client Service 119
- NetBackup Media Server 26, 40
 - and storage units 59, 228
 - network diagram of 34
 - selecting 57–58
- NetBackup Replication Director 17
- Network Attached Storage 26
- Network Attached Storage (data mover) 58
- network interface cards 261
- NEW_STREAM directive 90
- NIC cards and full duplex 261
- no-data Storage Checkpoint 95

O

- off-host backup 55, 57
 - and multiplexing 228
 - NAS 26
 - overview 25
 - prerequisites for 227
 - raw partition 228
 - type of disk (SCSI vs. IDE) 261
 - with data mover 43
- Open File Backup
 - disabling 79
 - license 43
- operating system
 - patches 35
- Oracle 36
- OST_FIM method 63
- overview of snapshot operations 287
- overwriting
 - raw partition restores 233

P

- page code 83 38, 228, 261
- pairresync command 268
- pairsplit (Hitachi) 202
- pairsplit (HP-XP) 206
- parameters for snapshots 65
- partitions
 - Windows 83
- patch for VxFS 36
- patches 35, 260
- pathname length 55, 76
- Perform block level incremental backups 52

- Perform snapshot backups 281
- performance
 - increasing tape 262
- peripherals (latest info on web) 43
- PFI_BLI_RESTORE file (for block-level restore) 235
- physical device (as raw partition) 124
- platform requirements 35
- platforms supported 45
- plex option (Snapshot Attribute for VSS) 69
- point-in-time snapshots 21
- policy
 - for NAS snapshot 120
 - how to select type of 53
 - storage unit 59
- Policy dialog 53, 84
- policy_name (on mover.conf file) 262
- primary vs. alternate client 27
- promotion
 - file 20, 235–236
- provider 18, 261
- Provider Type (for VSS) 68

Q

- query snapshot 247, 250, 265
- quiesce 288, 291

R

- RAID 5 131
- raw partition 87, 90
 - as snapshot source 77
 - backup 56
 - block vs. character device 228
 - defined 40
 - not supported with VxFS_Checkpoint 128
 - restore 232
 - fsck needed after vxfs restore 233
 - specifying for cache 124
- recovery procedure 264, 282
- Registry 82, 140
 - and FlashBackup 82
- remote snapshot (see alternate client backup) 27
- removing
 - clone 270
 - snapshots 264, 282
- replicated host 32
- replication
 - for alternate client backup 136
 - testing setup for alternate client backup 137

Replication Director 26, 35, 63, 263
 requirements for NetBackup 35
 restore 232
 and fsck 233
 block-level restore 235
 configurations 243
 FFR with vxvm or FlashSnap 102, 234, 236
 file promotion 20, 235–236
 from disk image 245–246
 from EMC_TimeFinder_Clone 181, 238
 from FlashBackup 232
 hardware-level 249
 logs 255, 257
 NAS_Snapshot 119
 Oracle files 235
 overwrite option 233
 raw partitions 232
 re. device file 233
 Restore everything to its original location 240, 242
 restrictions 36, 82
 resyncfromreplica option 248, 251
 resynchronization of mirror 132
 resynchronize
 disks 267
 Resynchronize mirror in background 69
 retention level for backup 101
 RMAN 40
 rollback 234, 238
 and clone creation 181, 238
 causes fractured (split) clone 211
 VSS and disk array credentials 158
 root
 specifying as snapshot source 77
 rotation 99
 of snapshots 100
 RSM Database 83, 140

S

SAN 36, 227
 defined 41
 Schedule tab
 Instant Recovery 103
 scripts
 running before/after backup 71
 SCSI E4 target descriptors 59
 SCSI Inquiry Page Code 83 38, 228, 261
 SCSI serialization 38, 228, 261
 SCSI vs. IDE 261
 serial numbers 38, 228, 261

serialization 38, 228, 261
 Shadow Copy Service (see VSS) 75, 139
 ShadowImage method 262
 shared disk group (VxVM) 124, 133
 SIZE_DATA_BUFFERS 262
 snap
 removing 266
 snapcachelist command 284
 snapctl 82
 driver log 256
 overview 281
 snaplist command 266, 283
 snapoff command 266, 285
 snapon command 282
 snapshot 21
 auto selection 59, 62, 74, 102
 back up to local storage 24
 backup to local storage 33
 configuration 60, 102
 controlling rotation 100
 copy-on-write vs mirror (how to choose) 24
 defined 41
 deleting 100
 disabling 79
 ID 247, 250, 265
 instant 108
 mirror
 defined 22
 mirror (creating) 134–135
 mirror (VxVM) 41, 131
 mirror access time 78
 naming format
 NAS snapshot 122
 on Solaris client
 troubleshooting 264
 overview 21
 pre-processing for 288
 removing 264–265, 282
 restoring from disk 246–247, 250
 rotation 99
 selecting method of 61
 source
 defined 41
 for symbolic link 77
 state file 48
 volume (creating) 134–135
 VxVM instant 65
 Snapshot Attribute (for VSS) 69

- Snapshot Client 21
 - access web info 43
 - installing 47
- Snapshot Client Options button 60
- Snapshot Options dialog 60–61, 74
- snapshot provider 18
- Snapshot Resources pane 100, 175
- Snapshot Volume
 - defined 41
- Snapshots only
 - on Schedule tab 102, 104
- snapstat command 285
- software option (Provider Type for VSS) 68
- Solaris
 - version requirements 35
- Solaris sd.conf
 - disk array methods 155
- space-optimized snapshots 65, 108, 131–132
 - snapshot methods supported 133
- split mirror backup (alternate client) 28
- SSSU for HP StorageWorks 184
- stack quiescing 290
- Standard policy
 - restoring 232
- state file 48
- status 13
 - and NetBackup Media Server off-host 229
- status code 156 260
- STD/BCV mirror pairs
 - EMC_TimeFinder_Mirror 180
- Storage Checkpoint 35, 82, 103
 - defined 41
 - determining disk usage 41
 - disk layout 94
- storage checkpoint 64
- storage devices 52
- storage unit 59
 - restrictions for data movers 59, 228
- storage_unit_name version of mover.conf file 261
- streams
 - allow multiple data 90
- support web site 43
- supported
 - data types 82
 - platforms 45
- switched fabric 39
- symbolic links 77
- symclone 181
- symmir 180

- symmir command 267
- symsnap 162
- synchronize disks 267
- Synchronize mirror before backup 69
- system option (Provider Type for VSS) 68
- system requirements for NetBackup 35
- System State 82, 140
- system-protected files 140
- system-protected files and FlashBackup 82

T

- tape
 - increasing performance 262
- Terminal Services Database 83, 140
- termination 264, 282
- terminology 38
- third-mirror (vxvm) 131
- third-party copy 26
 - and multiplexing 38, 228
 - and storage unit 59
 - defined 39
 - device configuration 57, 59
- third-party copy device 43
 - and storage units 59, 228
 - defined 42
- Third-Party Copy Device off-host backup 85
- TimeFinder method 262
- timestamp
 - incremental backup 73
- tpconfig 119
- troubleshooting 254
 - directories to create 254
 - disk array methods 209
- types of backups supported 86

U

- UFS file system 42, 63, 123
- UID 73
- umount command 248, 267
- UNC for Windows pathnames 121
- unified logging 255–256
- Universal Naming Convention (UNC) 121
- unmount
 - checkpoint 267
 - snapshot 248
- unquiesce 289
- UNSET directive 91–92
- UNSET_ALL directive 91–92

- user-directed
 - archives 232
 - backup 232

V

- VCMDB (Volume Configuration Management Database) 178
- vendors (latest info on) 43
- VERBOSE setting for logs 254
- Veritas Federated Mapping Services 42
- Veritas Volume Manager 65, 131
- Veritas Volume Manager cluster. *See* CVM
- Veritas Volume Replication 103
- virtual machine proxy 57
- VMware 57, 63
 - ALL_LOCAL_DRIVES 55, 104
- volume
 - defined 42
 - sets (VxVM) 95
- vshadow.exe 160
- VSP method 63
- VSS
 - disk array credentials and rollback 158
- VSS method 64, 139
- VSS snapshot method 103
- VVR 36, 103
- VVR method 32, 64, 137
 - preparing for alternate client backup 136
- vxassist 108–109, 131, 134–135, 266
- vxassist snapstart 107
- vxdbg command 107, 134, 136, 248, 251
- vxdbg list command 271
- vxdisk command 248
- VxFS clone
 - removing 266
- VxFS file system 35, 63, 82, 123
 - and AIX
 - Linux 45
 - patch for library routines 36
 - restoring 233
- VxFS multi-volume file system 36, 128, 131
- VxFS_Checkpoint method 64, 82, 103, 128, 235
- VxFS_Snapshot method 65, 130
- vxibc command 137
- vxmake 108
- VxMS 40, 42
- VxMS logging 257
- vxprint 107, 109, 266
- vxprint command 266

- vxrecover command 248
- vxsnap 110
- VxVM
 - and RAID 5 131
 - clone of disk group 270
 - instant snapshots 65, 109, 132
 - mirror 131
 - preparing for Instant Recovery 106
 - provider
 - with VSS 68
 - required version 62, 73, 130, 150
 - shared disk group 124, 133
 - Volume Manager 36, 131
 - volume name restriction 95, 106, 131
 - volume sets 95
- vxvm method 65, 103, 130–131, 237
- vxvol 107, 131, 134

W

- web access to recent Snapshot Client info 43
- whatrmver 201
- wildcards in Backup Selections list 56
- Windows
 - open file backup 43
 - OS partitions 83
 - System database files 83, 140
- Windows Shadow Copy Service 139