

# Symantec NetBackup in Highly Available Environments Administrator's Guide

for Windows, UNIX, and Linux

Release 7.6 Beta

# Symantec NetBackup in Highly Available Environments Administrator's Guide

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

Documentation version: 7.6 Beta

## Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

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

Symantec Corporation  
350 Ellis Street  
Mountain View, CA 94043  
<http://www.symantec.com>

DRAFT

# Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

## Contacting Technical Support

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

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

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

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

- Product release level

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

## Licensing and registration

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

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

## Customer service

Customer service information is available at the following URL:

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

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

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

## Support agreement resources

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

Asia-Pacific and Japan	<a href="mailto:customercare_apac@symantec.com">customercare_apac@symantec.com</a>
Europe, Middle-East, and Africa	<a href="mailto:semea@symantec.com">semea@symantec.com</a>
North America and Latin America	<a href="mailto:supportsolutions@symantec.com">supportsolutions@symantec.com</a>

DRAFT

# Contents

Technical Support .....	4
Chapter 1	About in this guide ..... 11
	What's in this guide ..... 11
	Documents related to NetBackup in highly available environments ..... 12
Chapter 2	NetBackup protection against single points of failure ..... 13
	About protecting against component failures ..... 13
	About protecting against network link failures ..... 15
	About protecting against storage device connection failures ..... 15
	About protecting against storage device failure ..... 16
	About protecting against media availability failures ..... 16
	About protecting against master server failures ..... 17
	About protecting against media server failures ..... 18
	About protecting against LAN client failures ..... 21
	About protecting against SAN client failures ..... 22
	About protecting against site failures ..... 22
	About protecting catalog in highly available environments ..... 23
	About installing media servers in clusters ..... 24
	About virtual storage units in a NetBackup cluster ..... 24
Chapter 3	About site disaster recovery with catalog backup and recovery ..... 31
	About catalog recovery ..... 31
	About full catalog recovery ..... 32
	Performing full catalog restore ..... 33
	Making the DR environment consistent after a full catalog restore ..... 35
	About partial catalog recovery ..... 35
	Performing partial catalog restore ..... 36

	Making the DR environment consistent after a partial catalog restore .....	37
	About disk recovery in DR domain .....	37
	Disk recovery in single-domain replication DR environment .....	38
	Auto Image Replication .....	38
	Disk recovery in cross-domain replication DR environment .....	38
Chapter 4	About site loss protection with catalog replication .....	41
	About auto image replication .....	41
	About conditions for support of replicated NetBackup catalogs .....	41
	About catalog synchronization .....	44
	About NetBackup catalog replication .....	44
	About multi-site single domain replication .....	44
	About multi-site cross domain replication .....	47
	About full catalog replication .....	49
	About partial catalog replication .....	52
Chapter 5	Deploying NetBackup master servers with full catalog replication .....	57
	About replication considerations .....	57
	About non-clustered NetBackup master server with catalog replication .....	58
	Installing and configuring non-clustered NetBackup master server with catalog replication .....	58
	About globally clustered NetBackup master servers with catalog replication .....	65
	Installing and configuring a globally clustered NetBackup master server with catalog replication .....	65
	Populating the server tables in the EMM .....	68
	Upgrading NetBackup in clustered replication configuration .....	69
	Failing over to secondary master server cluster .....	71
	Testing the NetBackup master server cluster in clustered replication environment .....	72

Chapter 6	Using NetBackup to perform backups and restores in a cluster .....	73
	About backups and restores with NetBackup in a cluster .....	73
	Performing user-directed backups with NetBackup in a cluster .....	73
	About restoring data in a cluster .....	74
	About supported NetBackup database agents and options in a cluster .....	76
	About backing up database files in a cluster .....	77
	About user backups .....	77
	About NetBackup client in a cluster .....	77
Index .....		79

DRAFT

DRAFT

# About in this guide

This chapter includes the following topics:

- [What's in this guide](#)
- [Documents related to NetBackup in highly available environments](#)

## What's in this guide

The *Symantec NetBackup in Highly Available Environments Administrator's Guide* discusses various methods for making NetBackup highly available and provides guidelines for protecting NetBackup against single point of failures.

This guide looks at the components of a data protection system based on NetBackup. It outlines different configurations and solutions for reducing the risk of failure within a particular site and recovering from the loss of the site.

You can use this guide as an aid to create a NetBackup site disaster recovery plan because it discusses catalog recovery and catalog replication process. However, this guide is not intended to provide a definitive disaster recovery plan for all NetBackup environments. Instead, you can use the information to develop site disaster recovery plans specific to your NetBackup environments.

This guide also provides guidelines for installing and upgrading NetBackup master servers. In addition, the guide details operating practices when catalogs are replicated between clustered or non-clustered NetBackup master servers.

The guide does not cover the details of clustering or replication technologies that are used. Please refer to your specific replication technology documentation for details on deploying and operating replication layers. For more information about NetBackup master server clustering, refer to the *Symantec NetBackup Clustered Master Server Administrator's Guide*.

See [“Documents related to NetBackup in highly available environments”](#) on page 12.

## Documents related to NetBackup in highly available environments

When you refer to the *Symantec NetBackup in Highly Available Environments Administrator's Guide* you may also want to refer to the following documents:

- Refer to *Symantec NetBackup Clustered Master Server Administrator's Guide* for information about clustering NetBackup.
- Refer to *Symantec NetBackup Installation Guide* for information about installing NetBackup.
- Refer to *Symantec NetBackup Administrator's Guide*, Volume I and Volume II for general information about NetBackup.

DRAFT

# NetBackup protection against single points of failure

This chapter includes the following topics:

- [About protecting against component failures](#)
- [About protecting against site failures](#)
- [About protecting catalog in highly available environments](#)
- [About installing media servers in clusters](#)

## About protecting against component failures

NetBackup comprises a number of different components, each of which has the potential to fail, and disrupt the backup or restore process.

[Table 2-1](#) lists the component level points of failure and the related protection method.

**Table 2-1** NetBackup protection against component failures

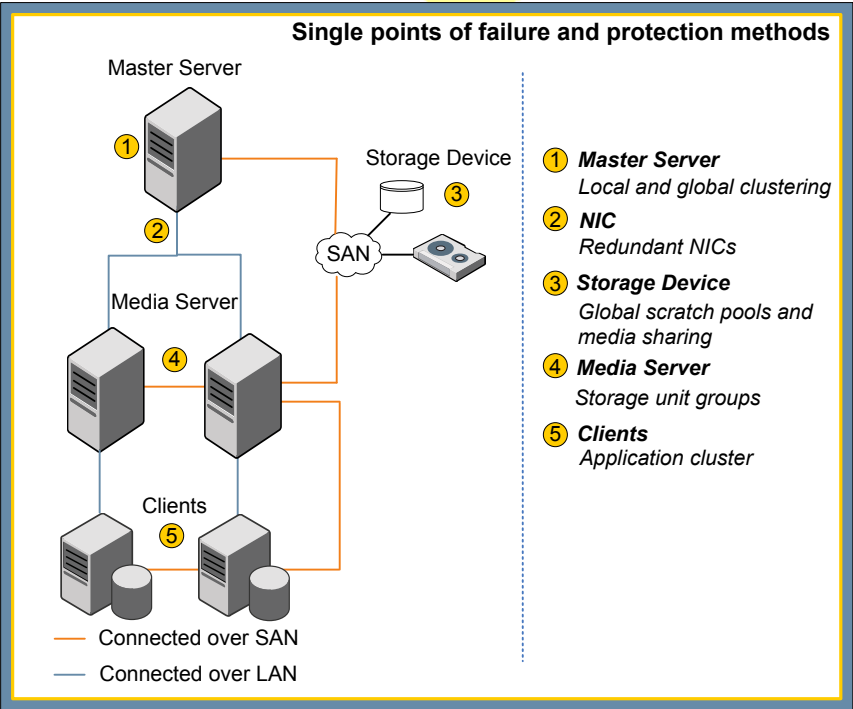
Point of failure	Protection method
Network links	See <a href="#">“About protecting against network link failures”</a> on page 15.
Storage device connections	See <a href="#">“About protecting against storage device connection failures”</a> on page 15.
Storage devices	See <a href="#">“About protecting against storage device failure”</a> on page 16.

**Table 2-1** NetBackup protection against component failures *(continued)*

Point of failure	Protection method
Media availability	See “About protecting against media availability failures” on page 16.
Master server	See “About protecting against master server failures” on page 17.
Media server	See “About protecting against media server failures” on page 18.
LAN client	See “About protecting against LAN client failures” on page 21.
SAN client	See “About protecting against SAN client failures” on page 22.

Figure 2-1 illustrates various NetBackup components and the single points of failure. The single points of failure can be eliminated at each component level either by making the component highly available or by deploying multiple components for redundancy.

**Figure 2-1** Single points of failure and protection methods



## About protecting against network link failures

The majority of backup traffic is transferred over network connections with 100 MB and 1Gbit speed which provide transfer rate of around 8 MB/sec and 65 MB/sec, respectively. To make network links highly available, deploy redundant network teaming. Due to cost considerations, network teaming is often restricted to backup servers and mission critical clients only. Non-mission critical clients have single network connections and the risk of connection failure (and the subsequent failure of the backup) is accepted.

## About protecting against storage device connection failures

Connections to storage devices and their controllers also represent single points of failure. In case of connection failure, the device cannot be used.

See [“About protecting against SAN connection failures”](#) on page 15.

See [“About protecting against robotic control connection failures”](#) on page 15.

### About protecting against SAN connection failures

SAN connections generally exist between the backup servers and the backup storage; although the NetBackup SAN client also supports SAN connections from clients to media servers. In all cases, to protect NetBackup against SAN connection failure, SANs should be configured to provide redundant connections between the source and the target components.

Most SAN-attached disk arrays have redundant SAN connections and support dynamic multi-pathing (DMP) software. This redundancy ensures that the connection to the storage is maintained even if one path fails. In many cases, DMP software also load balances traffic across SAN connections to improve the data transfer rates to and from the disk storage.

Many SAN-attached tape devices also offer two connections for redundancy, and thus they appear to servers as two separate devices. Multi-path selection is not dynamic. NetBackup selects the first available path it finds and always uses that path. The second device path is only used if the first path is broken.

### About protecting against robotic control connection failures

In tape-based backup environments, the robotic control connections can be single points of failure. The inability to send instructions to the tape library prevents backup and restore operations, even if the tape drives are available.

Some tape libraries, such as Sun STK ACSLS or Quantum ATM, use a dedicated control software that runs on a server that is independent of the library. Such control servers can be clustered. The media servers send requests to the control

server, which handles the movement of tapes between slots and drives in the library.

Other tape libraries depend on a direct device connection from the NetBackup master server for control instructions to the library. If this device connection is lost, the tape library cannot be used. SAN-attached tape libraries support multiple connections to the robotic control for redundancy. You can configure these connections to provide protection against server failure. For example, you can configure one path to each node of a clustered master server. You must ensure that the paths are not active at the same time. If both paths are active, conflicting instructions can be issued, which could result in backup failure or data loss.

## About protecting against storage device failure

Whether they are tapes or disks, when storage devices fail they are considered to be single points of failure. To protect against storage device failures, you should have multiple devices as backup targets.

A media server with access to only one tape drive cannot complete backups to tape if that tape drive goes down. To protect NetBackup against such failures, configure the media servers to access at least two tape drives. Use SAN-attached tape drives, which can be shared between media servers. This sharing ensures that the tape drives are accessible without needing large numbers of redundant devices. Typically, one or two redundant drives provide for resilience and allow restore operations to occur while backups are in progress. For example, if you configure four media servers to share five tape drives, backups can still happen even if one drive goes down. The backup may take longer, but it completes and your data remains safe. If media servers run backups at different times, the ratio of tape drives to servers may be even lower without risking backup failure.

AdvancedDisk disk pools can be created on individual media servers to protect against the failure of a single disk device.

## About protecting against media availability failures

In tape-based backup solutions, failures can occur if no suitable tape media is available for use by a backup job. With NetBackup, risk of such failures can be reduced through global scratch pools and media sharing.

[Table 2-2](#) discusses the methods of protection against media availability failures.

**Table 2-2** NetBackup protection against media availability failures

Protection method	Description
Global scratch pools	<p>For all the backup jobs and duplication jobs that are written to tapes, use the tapes that are in a specific media pool with the same retention criteria as the backed up data. If no suitable tapes are available, the backup fails.</p> <p>A global scratch pool is a NetBackup media pool that holds unassigned tapes that can be automatically re-assigned to a specific media pool on demand. For instance, a backup or a duplication job runs and no suitable tapes are available in the media pool specified by the job. Then an unassigned tape is transferred from the global scratch pool to the specified media pool and is used for the backup job. When this tape expires, it is automatically returned to the global scratch pool for re-use.</p> <p>Using a global scratch pool ensures that all unassigned tapes are available for use by any backup job, irrespective of the media pool specified by the job.</p>
Media sharing	<p>Media sharing allows multiple media servers to use partially full tapes until they are full. It ensures the most efficient use of tape. Only one media server at a time can write to a tape. When that tape is not in use, a different media server that requires a tape from that media pool can use it.</p> <p>To enable media sharing, set the <b>Volume Pool</b> properties to use the <b>Maximum number of partially full media</b> property. This property restricts the number of partially full tapes in a media pool. Until all tapes are full, empty tapes cannot be assigned to the pool. Until one tape is full, another empty tape cannot be assigned to the pool.</p>

## About protecting against master server failures

A single master server for each NetBackup domain controls all the backup activity within the domain. Thus, the master server represents the most obvious single point of failure in the data protection environment. Without the master server, backups and restores are not possible. To protect NetBackup against such failures, the master servers must be highly available. NetBackup supports clustering of the master servers using the following clustering technologies:

- Microsoft Cluster Server
- Veritas Cluster Server Windows 2003, 2008, 2008R2
- Veritas Cluster Server on UNIX/Linux
- HP Service Guard

■ HACMP for AIX

More information about installing and configuring NetBackup on these cluster technologies is available in the *Symantec NetBackup Clustered Master Server Administrator's Guide*.

## About protecting against media server failures

Although media servers can be configured with redundant network and SAN connections, the servers themselves remain single points of failure. Methods of protecting NetBackup against media server failures may vary depending on the type of media servers that you use.

[Table 2-3](#) lists the different types of media servers and the protection method.

**Table 2-3** Type of media servers and protection method

Type of media server	Description
Dedicated media servers	Run only the media server software and exclusively back up data from other systems.  See <a href="#">“About protecting against dedicated media server failures”</a> on page 18.
Non-dedicated media servers	Run other applications also that require backing up. Also back up data from other systems.  See <a href="#">“About protecting against non-dedicated media servers failures”</a> on page 19.
SAN media servers	Run other applications also that require backing up. Do not back up data from other systems.  See <a href="#">“About protecting against SAN media server failures”</a> on page 20.

## About protecting against dedicated media server failures

Storage unit groups can be used to protect NetBackup against the failure of a single media server. Storage unit groups can also be used for load balancing across multiple media servers to ensure optimal backup and restore performance.

[Table 2-4](#) discusses the different modes in which you can configure the storage unit groups.

**Table 2-4** Modes for configuring storage unit groups

Mode	Description
Failover	In the failover mode, the first storage unit is always used, unless the media server is down. Excess jobs are queued rather than being directed to the next storage unit. The failover mode functions similarly to what would be seen if two media servers were configured as an active or a passive cluster.
Prioritized	In the prioritized mode, the first available storage unit in the list is used. In this mode, jobs that exceed the total number the storage unit can handle, are directed to the next storage unit in the list. If the media server is down, all backups are directed to the next storage unit.
Round robin	In the round robin mode, different storage units from the list are used in a cycle for each job. If each storage unit is on a different media server, this acts as a load balancing mechanism.
Load balanced	The load balance mode only works with Flexible Disk and Media Manager storage unit types. In the load balance mode, NetBackup carries out checks on activity and resources available on each media. The check is carried out before the backup are directed to the media with the lightest load.

As a best practice, when using prioritized and failover groups to configure two storage unit groups, use two media servers, as follows:

- Configure each media server to have a single storage unit. For example, so Node A has STU A and Node B has STU B.
- Configure two storage unit groups with the storage units in a specific order in each one. In this example, SUG AB contains STU A, followed by STU B. SUG AB contains STU B followed by STU A.
- Backup policies are then evenly shared between SUG AB and SUG BA.

During operation, the backup traffic is normally shared between the two nodes, but if one node fails, all backups automatically go to the other node.

## About protecting against non-dedicated media servers failures

Storage unit groups can also be used to protect against the failure of non-dedicated media servers. However such use does not protect other applications running of a given media server from the failure of that media server. In some cases non-dedicated media servers may form part of cluster supporting other applications. These applications can be protected using virtual storage units.

See [“About virtual storage units in a NetBackup cluster”](#) on page 24.

## About protecting against SAN media server failures

Unlike regular media servers, SAN media servers only protect themselves. A SAN media server connects directly to the backup storage in the same way as a regular media server. But it does not receive data from other client systems over a network or SAN link.

SAN media servers are usually deployed on the servers that support large, mission-critical applications, which are often clustered. While the application may be clustered, you do not need to cluster the SAN media server itself. Instead, install the SAN media server software on each member node of the cluster and create application cluster definitions in the NetBackup EMM database for each virtual name the cluster uses. Then create a storage unit using the virtual name of the cluster as the media server. The associated application with a given virtual name use the storage unit that is associated with the same virtual name for backups.

See [“About installing media servers in clusters”](#) on page 24.

## Restoring tape backups using an alternative media server

Generally, while restoring files, NetBackup expects to use the same media server and client that it used for the original backup. However, for disaster recovery, you use a different media server to restore the backup to a different client. The media servers and clients at the disaster recovery site are likely to have different names from those at the primary site.

NetBackup lets you configure failover restore media servers to handle restores in the event that the original media server is unavailable.

To configure failover restore media servers:

- On Windows master server, you can configure failover restore media servers using the **NetBackup Administration Console**.  
Go to **Host Properties > Master Server > Restore Failover**.
- On UNIX and Linux master servers, you must create `FAILOVER_RESTORE_MEDIA_SERVER` entry in the `bp.conf` file.

## Restoring disk backups using an alternative media server.

NetBackup can share disk storage pools between multiple media servers. During restore, by default, NetBackup balances the job load and automatically directs the restore to the least busy media server rather than the one that made the backup. However, this process can cause problems if the media server selected to perform

the restore is licensed as a SAN media server or does not have network access to the client which requires a restore.

There are three options available if you encounter this problem:

- Configure the force restore media server setting as follows:
  - On UNIX and Linux master servers, you create `FORCE_RESTORE_MEDIA_SERVER` entry in the `bp.conf` file.
  - On Windows master server, you can define this setting in the **NetBackup Administration Console**.  
Go to **Host Properties > Master Server**.  
This setting works on a per-server basis. It lets you specify a media server for restore operations based on the media server that is used to make the backup. To ensure that the same media server is used to make the backup and the restore, specify the same name for the backup and restore server.
- Create the touch file `USE_BACKUP_MEDIA_SERVER_FOR_RESTORE`, as follows:
  - On UNIX and Linux master server, create the file in `/usr/opensv/netbackup/db/config`
  - On Windows master server, create the file in `<install path>\veritas\netbackup\db\config`.  
`USE_BACKUP_MEDIA_SERVER_FOR_RESTORE` is a global setting and always forces restore to the server that did the backup.

---

**Note:** When the `USE_BACKUP_MEDIA_SERVER_FOR_RESTORE` touch file is created, all `FAILOVER_RESTORE_MEDIA_SERVER` and `FORCE_RESTORE_MEDIA_SERVER` settings are ignored.

---

- Run the restore from the command line using the `bprestore -disk_media_server` command. This setting works on a per job level. It also lets you specify the media server that is required for the specific restore job. Unlike the other two options, this setting is dynamic and can be applied when needed.

## About protecting against LAN client failures

The NetBackup client package (including the application agents) is not cluster aware and must be installed separately on each node of a cluster that is being protected as a NetBackup client. When backing up clustered applications specify the virtual server name associated with the application as the client name in the

backup policy. This will ensure that the correct node of the cluster is selected during the backup operation.

## About protecting against SAN client failures

The SAN client, like the SAN media server, does not send backup traffic over the network to the media server. However unlike SAN media servers, which send backup data directly to the storage devices, SAN clients send backup data over a SAN connection to a remote media server.

SAN clients are often used to protect clustered applications. To protect NetBackup against SAN client failures when used in this way, configure the SAN client as application clusters in EMM. This configuration also ensures that the media server controlling the backup always opens a fiber transport connection to the active node of the cluster when a backup is initiated.

## About protecting against site failures

Local clustering provides local failover for each site. However, these configurations do not provide protection against large-scale disasters such as major floods, hurricanes, and earthquakes that cause outages for an entire region. The entire cluster can get affected by such an outage. In such situations, global clustering or wide area clustering ensures data availability by migrating applications to the remote clusters that are located considerable distances apart.

Global cluster architecture supports deployment of two or more datacenters, clusters, and subnets that are separated by a larger distance. A global cluster with replicated master server cluster can monitor and manage the replication jobs and clusters at each site. In case of a site outage, it controls the shift of replication roles to the secondary site. It brings up the critical applications and redirects client traffic, from one cluster to the other.

Auto image replication is a NetBackup feature which allows individual disk based backups to be replicated between NetBackup domains. Because the backups are automatically recorded in the NetBackup catalog of the target domain there is not need for catalog replication of complex catalog recovery procedures when using auto image replication. For more information, refer to the *Symantec NetBackup Administrator's Guide, Volume I*.

# About protecting catalog in highly available environments

The NetBackup catalog contains information about both existing backups and the backup policy, including what gets backed up when and to where and how long the backup is kept for. As such the catalog is a single point of failure and needs to be protected. Using the RAID storage provides some protection against storage failure. Replication can also protect against storage failure and site loss. Regular backups of the catalog can protect against corruption and accidental data loss.

See [Table 2-5](#) on page 23. discusses the various methods for protecting NetBackup catalogs.

**Table 2-5** NetBackup catalog protection in highly available environments

Protection Method	Description
Catalog backups	<p>The catalog backup protects the NetBackup catalog on the master server against both hardware failure and data corruption and catalog backups should be made on a regular basis, ideally at least daily. The catalog backup is policy-based so it has all of the scheduling flexibility of a regular backup policy. As the policy allows for incremental backups, catalog backup times for large catalogs can be significantly reduced. However it should be noted that recovery from incremental backups can take longer due to the need to restore.</p> <p>Catalog backups written to tape use media from the Catalog Backup volume pool only.</p> <p>For more information, refer to the <i>Symantec NetBackup Administrator's Guide, Volume I</i>.</p>

Table 2-5

NetBackup catalog protection in highly available environments

(continued)

Protection Method	Description
Catalog replication	<p>Catalog replication is the process of creating and managing duplicate versions of a catalog database. Catalog replication copies a database and synchronizes a set of replicas so that the changes that are made to one replica are reflected in all the others.</p> <p>Replicating the catalog to a standby master server at the disaster recovery or secondary site ensures rapid catalog recovery at the disaster recovery site. Continuous replication ensures that the catalog is as up to date as the replication link allows.</p> <p><b>Note:</b> Replication does not protect against catalog corruption or accidentally deleting or expiring images. You must make regular scheduled catalog backups.</p> <p>See “<a href="#">About NetBackup catalog replication</a>” on page 44.</p> <p>See “<a href="#">About catalog recovery</a>” on page 31.</p>

## About installing media servers in clusters

In cluster configurations, you can install NetBackup media servers and SAN media servers as non-failover applications on each node of the cluster, if NetBackup master server is also a media server. Configuring in this way allows all nodes of the cluster to function as media servers, either in conjunction with storage unit groups or virtual storage units that fail over with other applications running on the same cluster.

Configure each media server as you would in a non-clustered environment. For more information, refer to the *Symantec NetBackup Installation Guide*.

See “[About virtual storage units in a NetBackup cluster](#)” on page 24.

## About virtual storage units in a NetBackup cluster

In a typical NetBackup environment, the storage unit is associated with a NetBackup entity, such as the master server or media server. In the configuration that is described in this topic, however, a storage unit is associated with the virtual name of the application in the cluster environment. The storage unit for that application is not associated with any specific node in the cluster and can failover to any node in the cluster. This lack of association ensures that the storage unit follows along with the application, wherever the application is running in the

cluster. The NetBackup master server uses the virtual name of the application as the media server name when it backs up the clustered application.

This configuration is supported for Veritas Cluster Server (VCS), Microsoft Cluster Server (MSCS), SunCluster, Service Guard, and HACMP.

## Prerequisites when configuring a NetBackup media server in a cluster with virtual storage units

The following requirements must be met when you configure a standalone NetBackup media server in a cluster with virtual storage units:

- Your hardware must support the clustering technology and NetBackup. For a list of NetBackup supported storage devices, visit the following Web site. Also, refer to the documentation that is included with your clustering software.  
<http://entsupport.symantec.com>
- Verify that storage devices are properly installed and configured to work in the cluster environment. Refer to the information on how to configure storage devices in the *NetBackup Installation Guide*.
- An IP address and host name (virtual name) must be associated with each clustered application.
- Verify that your cluster (VCS, MSCS, SunCluster, Service Guard, or HACMP) is correctly installed and configured.
- Make the NetBackup master server the EMM server and the robotic control host. (You may not be able to change the robotic control host in all configurations since not all robot types have a robot control option.)
- Verify that you have the NetBackup Enterprise Server 7.6 installation program and a valid license key. You must also have a valid license key for the NetBackup Shared Storage option, if you perform backups to tape.

## About installing NetBackup media software in a cluster

Install the NetBackup media server software on each node in the cluster. See the *NetBackup Installation Guide* for details.

---

**Note:** When you configure NetBackup as described in this topic, do not configure the NetBackup media server as a failover application. The media server must be treated as a stand-alone application on each node in the cluster.

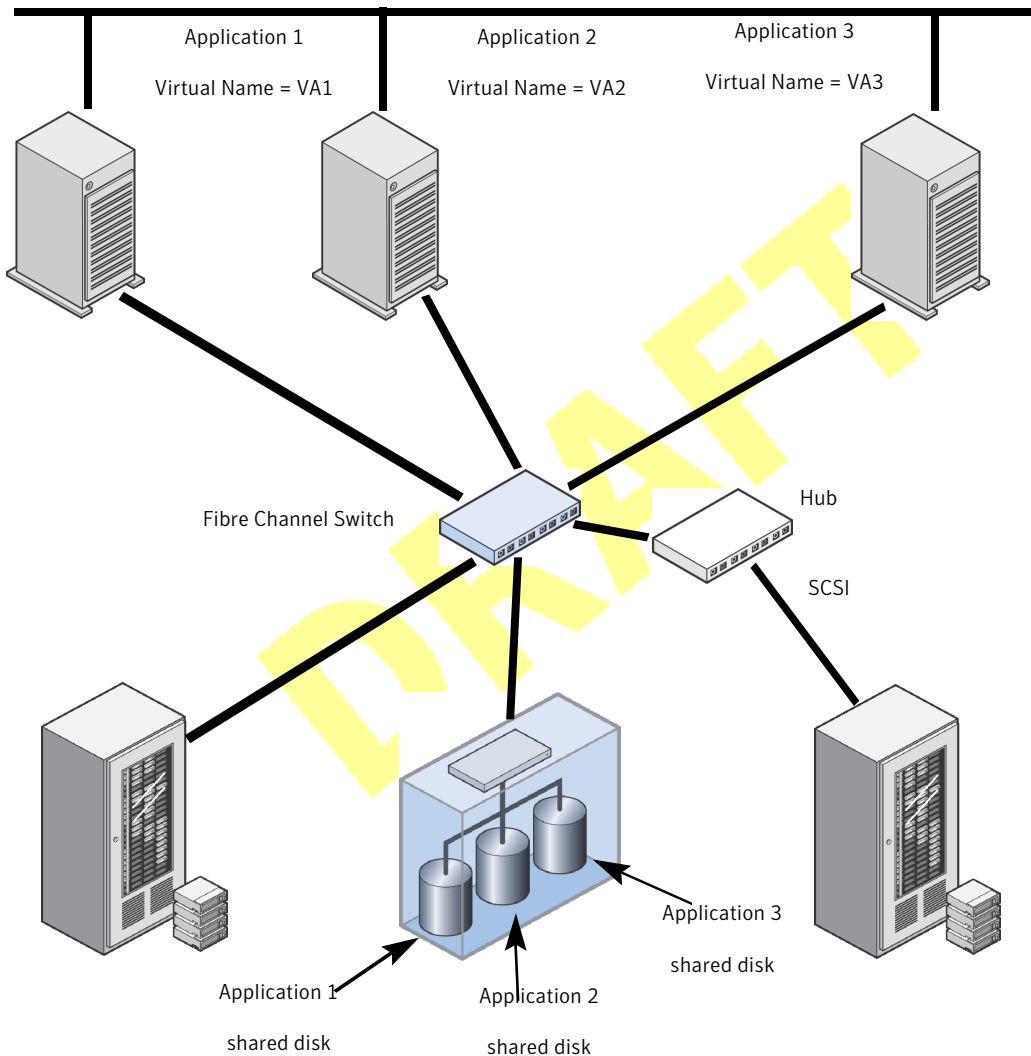
---

## About configuring virtual storage units in a NetBackup cluster

Storage units are configured differently depending on whether the storage is disk-based or tape-based.

Figure 2-2 shows NetBackup as installed on local media servers.

**Figure 2-2** NetBackup as installed on local media servers



## Adding tape storage unit

The following procedure describes how to add a tape storage unit.

### To add tape storage unit

- 1 Ensure NetBackup media server and the Shared Storage Option are installed and configured on all nodes in the cluster.
- 2 On the master server, make the following changes to the **Servers** list.

- Add the hostname for each node of the cluster.
- Add the virtual name for each highly available application.

For a UNIX/Linux environment, update the `bp.conf` file. For Windows, edit the **Servers** list in the host properties for the server.

- 3 On the master server, run the following command to add a host with the virtual name of the application:

```
nbsmmcmd -addhost -machinename <cluster_name> -machinetype
app_cluster
```

where `<virtual_name>` is the virtual name of the application.

- 4 On the master server, run the following command to associate the host with each node where the application runs in the cluster:

```
nbsmmcmd -updatehost -add_server_to_app_cluster -machinename
<node_name> -machinetype media -clustername <virtual_name>
-netbackupversion 7.6 -masterserver <master_name>
```

where the `<node_name>` is a node where the application can run.

`<virtual_name>` is the virtual name of the application. `<master_name>` is the name of the NetBackup master server.

- 5 Repeat step 4 for each node where the application can run.
- 6 On each node where the media server is installed, make the following changes to the **Servers** list.
  - Add the hostname for each node of the cluster.
  - Add the virtual name for each highly available application.

- 7 Use the NetBackup Administration Console to create a storage unit for each application. Specify the virtual name of the application as media server name for the storage unit.
- 8 To back up each application, create a policy that specifies the storage unit you created in step 7 for that application.

Allocate specific tape devices for this configuration so that you do not stress storage units jobs. You want to create storage units for each application within the cluster. Or, you may want to create storage unit groups which many applications can use.

## Adding disk storage units

Symantec recommends that you use the NetBackup Administration Console to add disk storage units. (You can also perform this task from the command line.)

### To add a disk storage unit

- 1 On the master server, make the following changes to the **Servers** list.

- Add the hostname for each node of the cluster.
- Add the virtual name for each highly available application.

For a UNIX/Linux environment, update the `bp.conf` file. For Windows, edit the **Servers** list in the host properties for the server.

- 2 On the master server, run the following command to add a host with the virtual name of the application:

```
nbemmcmd -addhost -machinename <virtual_name> -machinetype  
app_cluster
```

where `<virtual_name>` is the virtual name of the application.

- 3 On the master server, run the following command to associate the host with each node where the application runs in the cluster:

```
nbemmcmd -updatehost -add_server_to_app_cluster -machinename  
<node_name> -machinetype media -clustername <virtual_name>  
-netbackupversion 7.6 -masterserver <master_name>
```

where `<node_name>` is a node where the application can run.

`<virtual_name>` is the virtual name of the application.

`<master_name>` is the name of the NetBackup master server.

- 4 Repeat step 3 for each node where the application can run.

- 5 On each node where the media server is installed, make the following changes to the **Servers** list.
  - Add the hostname for each node of the cluster.
  - Add the virtual name for each highly available application.
- 6 Use the NetBackup Administration Console to create a disk storage unit for each application. Specify the virtual name of the application as media server name for the storage unit.
- 7 To back up each application, create a policy that specifies the storage unit you created in step 6 for that application.

## Upgrading a NetBackup media server in a cluster with virtual storage units

The following procedure describes how to upgrade a standalone NetBackup 6.0 or higher media server in a cluster with virtual storage units.

### To upgrade a NetBackup media server in a cluster with virtual storage units

- 1 On each node where NetBackup media server is installed, install the software on the local disk as you would in a non-cluster environment.  
 Follow the instructions on how to upgrade as described in the *Symantec NetBackup Installation Guide*.
- 2 On the master server, run the following command:
 

```
nbeemcmd -updatehost -machinename <virtual_name> -machinetype  
app_cluster -netbackupversion 7.6
```

where <virtual\_name> is the virtual name of the application.
- 3 Repeat step 2 for each virtual application that you have in your cluster environment.

DRAFT

# About site disaster recovery with catalog backup and recovery

This chapter includes the following topics:

- [About catalog recovery](#)
- [About disk recovery in DR domain](#)

## About catalog recovery

A major problem users encounter during site disaster recovery is that the disaster recovery (DR) site is not a mirror image of the production site. To perform DR operations you need a copy of the NetBackup catalog from the production master server. The NetBackup catalog backup and recovery process is primarily intended for recovering from catalog storage or master server failure rather than site loss. The default situation is that NetBackup restores the complete catalog including the EMM database. The EMM database includes details of the media servers, backup devices, and storage units. Master server use this information to direct backups and restores. Master servers also use this information to interrogate the media servers, to establish the status of the backup devices. In a DR environment which does not contain these media servers, the performance of the master server can be affected. Also, the ability to carry out restore operations can be affected, as polling operations fail to connect and time out.

Use the following approaches to recover the NetBackup environment at a DR site where the arrangement of media servers and clients is different from the main production site. Both approaches have advantages and disadvantages.

- In the full catalog recovery approach the whole catalog is recovered and then unwanted configuration elements can be removed or disabled.  
See [“About full catalog recovery”](#) on page 32.
- In the partial catalog recovery the EMM and the BMR databases are not restored.  
See [“About partial catalog recovery”](#) on page 35.

The most appropriate method for recovery can be determined by the nature of the DR facility and how similar it is to the production facility.

When creating your disaster recovery plan, ensure that it is in line with the approaches discussed in the following sections:

- See [“Planning a cross domain replication disaster recovery domain”](#) on page 48.
- See [“Performing full catalog restore”](#) on page 33.
- See [“Performing partial catalog restore”](#) on page 36.

## About full catalog recovery

Full catalog recovery is primarily used to recover the catalog if the data is corrupted or storage is lost at the production site. Full catalog recovery is recommended for single domain configurations. Full catalog recovery is used if the DR site has the same number of media servers with the same names as those used at the production site.

Full catalog recovery has the following advantages over partial catalog recovery:

- It restores the relational database components, which include the storage unit definitions, media assignment, and history.
- It retains the tape information from the primary site including the media pool and other assignment information.
- It restores the BMR data.
- It enables backups to be run at the DR site using the same policies and tapes that are used at the production site.

With full catalog recovery, there are the following limitations:

- When you recover the relational database components, the device configuration and the server configuration set up at the DR site before recovery is lost. You must set it again after recovery. The information that exists in the relational database about production servers and devices may not exist at the DR site. To ensure smooth operation in the DR environment, these server entries must be disabled and the devices associated with them should be removed.

- Full catalog recovery overwrites the device configuration and the server configuration in the relational database. You must rediscover the DR domain server and device configuration after the catalog is restored.

## Performing full catalog restore

With full catalog recovery the complete catalog backup is recovered to the DR master server. The media servers that do not exist in the DR environment are deactivated to avoid unnecessary pooling. All device records are removed because the device configuration at the DR site can be different to the production site. Device discovery is run to update the EMM database. You must perform the following procedure before restores can be started. Also, document the procedure in your DR plan.

### To prepare for full catalog restore

- 1 On UNIX and Linux master servers, create copies of the `bp.conf` and `vm.conf` files.
- 2 Run the `bprecover` command to recover the entire catalog.

---

**Note:** The DR master server must have the same name and topology as the production master server. If the production master server is a cluster then the DR master server must also be a cluster. The number of member nodes and the names of the nodes can be different.

---

---

**Note:** If a catalog backup that was created on a separate media server is used, then a media server with the same name is required for the catalog recovery.

---

- 3 Deactivate all the backup policies to prevent backups from starting automatically.
  - You can do this manually using the **NetBackup Administration Console**
  - Or run the `bppllist <policy> -set -inactive` CLI.
- 4 Shut down NetBackup.
- 5 On UNIX and Linux master servers, replace the `bp.conf` and the `vm.conf` files that were restored from the catalog backup with the copies created in step 1.
- 6 Start the NetBackup Relational Database Manager, NetBackup PBX, and EMM services on the new master server.
  - On UNIX and Linux master servers, run the following commands:
    - `/usr/openv/netbackup/bin/nbdbms_start_stop start`

- `start /opt/VRTSpx/bin/pbx_exchange`
- `/usr/opensv/netbackup/bin/nbemmm`

---

**Note:** For NetBackup versions 6.5.3 and higher, run the `nbemmm -maintenance` command.

---

- On Windows master servers, start the following Windows services:
  - NetBackup Relational Database Manager
  - Symantec Private Branch Exchange
  - NetBackup Enterprise Media Manager

---

**Note:** The PBX process may already be running because the NetBackup commands do not stop and start PBX.

---

For more information about NetBackup Relational Database Manager service, see the *Symantec NetBackup Troubleshooting Guide*.

- 7 Deactivate the media servers that are not part of the DR environment. Run the following command:

```
nbemmmcmd -updatehost -machinename <Media Server> -machinestateop  
set_admin_pause -machinetype media -masterserver <Master Server>
```

- 8 Delete all the tape devices from the EMM database. Run the following command:

```
nbemmmcmd -deleteallddevices -allrecords
```

- 9 Restart NetBackup.
- 10 Using the **Device Configuration** wizard create the new tape drive and library configuration.
- 11 If bar code masking rules were used at step 6, ensure that the same rules are set here. If necessary, add them.
- 12 Using the **NetBackup Administration Console**, verify if all the recovery media are set to non-robotic.
- 13
  - If some recovery media still need to be set to non-robotic, do the following:
    - Select the robotic media, right-click and select **Move**.
    - Change the robot field to **Standalone**.

- Click **OK** to save the changes.

**14** Once all the recovery media are set to non-robotic, in the **Inventory all the tapelibraries** field ensure that the media are identified in the correct library.

You can now start restore and recovery operations of the client data that is backed up at the production datacenter .

## Making the DR environment consistent after a full catalog restore

In the event of a major incident at the production site, operate from the DR site for some time after the basic recovery is completed. The following additional tasks may be optionally carried out once the DR environment is operational to make the DR environment consistent.

### To make the DR environment consistent

- 1 Modify the backup policies, including the **catalog backup** policy, to use the storage units available at the DR site **and enable them**.
- 2 Delete the backup policies that **are no longer required**.
- 3 Delete the storage units that are associated with the media servers and are not part of the DR **environment**.
- 4 Modify any Storage Lifecycle Policies that use storage units that you have deleted.

## About partial catalog recovery

Partial catalog recovery is recommended for multi-domain configurations. Partial catalog recovery is used for DR sites where the server layout is different from the production site with fewer media servers, different library types, etc. Partial catalog recovery is a variation of the Recovery without import method. It is subjected to many of the same constraints. For more information, go to the following link:

Partial catalog recovery recovers only the flat file components and not the relational database. Thus, the details of the existing infrastructure (servers, devices etc.) at the DR site is not lost during the recovery process. It also means that the media server information that is associated with the backups is not recovered. The media server must be manually added to the database and is unassigned. Ensure that the media server are placed in a pool where they cannot get accidentally overwritten.

Partial catalog recovery has the following advantages over full catalog recovery:

- No elements of the configuration need to be removed or rediscovered. The recovery process does not affect the general configuration of the DR environment.
- It does not affect the server topology. The master server topology at the DR site does not need to reflect the topology at the production site. Thus, a catalog backup from a clustered master server can be restored to a standalone master server at the DR site.
- The DR site can be a production site, provided the client names, backup policy names, and tape label ranges used in the two environments are unique. Also, it must be possible to do a partial recovery to another production backup domain.

With partial catalog recovery, you cannot recover the tape information from the primary site at the DR site. Ensure that the tapes are **not** accidentally overwritten. These tapes must not be easily used for backups at the DR site.

## Performing partial catalog restore

With partial catalog approach, it is assumed that restore operations do not need tapes to be assigned or located in specific media pools. It is also assumed that a tape exists in EMM and NetBackup can mount and read the tape for restoring. The following steps must be carried out before restores can be started:

### To prepare for partial catalog restore

- 1 On UNIX and Linux master servers create copies of the `bp.conf` and `vm.conf` files.
- 2 Recover only the NetBackup catalog image and configuration files.
  - When using the **NetBackup Administration Console**, select the **Partial catalog recovery** option when prompted.
  - Or run the `bprecover -wizard` command.

---

**Note:** The DR master server must have the same name as the production master server.

---

---

**Note:** If a catalog backup that was created on a separate media server is used, a media server with the same name is required for the catalog recovery.

---

- 3 Run the `cat_export -all -staging` to export the metadata from the replicated relational database backup.

- 4 Run the command `cat_import -all` to import the exported metadata into the active relational database. Alternatively, set the parameter `LIST_FS_IMAGE_HEADERS` to `YES` in the `bp.conf` file or the registry depending on the master server platform. This will cause the next catalog cleanup job to automatically import the exported metadata.
- 5 Deactivate all the backup policies to prevent backups from starting automatically.
  - You can do this manually using the **NetBackup Administration Console**.
  - Or run the `bppllist <policy> -set -inactive` CLI.
- 6 Shut down NetBackup.
- 7 On UNIX and Linux master servers, replace the `bp.conf` and the `vm.conf` files that were restored from the catalog backup with the copies created in step 1
- 8 Start NetBackup.
- 9 Inventory all the tape libraries to ensure that the tapes are added to the non-scratch media pool. This pool prevents tapes from being accidentally overwritten by active backup policies at a later time.

You can now start restore and recovery operations of client data that is backed up at the production datacenter.

## Making the DR environment consistent after a partial catalog restore

In the event of a major incident at the production site, operate from the DR site for some time after the basic recovery is completed. The following additional tasks may be optionally carried out once the DR environment is operational to make the DR environment consistent.

### To make the DR environment consistent

- 1 Modify and enable backup policies, and the catalog backup policy, that is required at the DR site.
- 2 Delete the policies that are no longer required.

## About disk recovery in DR domain

With introduction of OpenStorage and other AdvancedDisk types, deduplication disk as a backup storage medium is preferred over tape storage. Using disk storage you can replicate the contents of a disk device to another disk device in a secondary location. This replication eliminates the need to transport the physical backup media to a disaster recovery site.

## Disk recovery in single-domain replication DR environment

You can use the storage lifecycle policies, to optimize replication of deduplicating disks when duplicating backups within the same NetBackup domain. This is an efficient way to create duplicate copies of backup images at a disaster recovery site, which is controlled by the same master server as the production site. However, optimized deduplication is effective only for single-domain replication.

## Auto Image Replication

Auto Image Replication extends the concept of duplicating backups to separate domains, allowing individual backup copies to be sent to a DR domain. As backup copies created using auto image replication are automatically cataloged in the DR domain there is not need for additional recovery steps within the DR domain. Refer to *Symantec NetBackup Administrator's Guide, Volume I* for more information about auto image replication.

## Disk recovery in cross-domain replication DR environment

If the disk technology being used does not support Auto Image Replication, an alternative approach is simply to replicate the entire storage and then use a combination of catalog recovery and the `nbcatsync` utility to populate the catalog at the disaster recovery location. You can use the `nbcatsync` utility for cross-domain replication. Using the `nbcatsync` utility, you can replicate disk storage between the production domain and the disaster recovery domain without importing the entire contents of the replicated disk storage.

The `nbcatsync` utility facilitates replication even if disk media IDs recorded at the EMM database and at the metadata component of the image database are different. The `nbcatsync` utility aligns the disk media IDs in the image database metadata with the media IDs in the disaster recovery domain's EMM database. The regular backups and catalog backups that are made at the production site are written to the replicating disk storage. The catalog backup's disaster recovery file is sent to the disaster recovery domain.

The `nbcatsync` utility is supported on all master server platforms. You can use it with all Advanced Disk types supported by NetBackup.

**To recover disk in a cross-domain replication environment in the event of a disaster, perform the following steps on the DR domain's master server:**

- 1** Align the disk media ID information in the catalog backup's DR file with the disk media ID information in the DR domain's EMM database. For this, run the following command:

```
nbcatsync -sync_dr_file <DR file name>
```

- 2** Perform a partial catalog recovery from the replicated catalog backup by running the command,

```
bprecover -wizard
```

- 3** Run the command `cat_export -all -staging` to export the metadata from the replicated relational database backup.
- 4** Run the command `cat_import -all` to import the exported metadata into the active relational database.
- 5** Align the disk media IDs associated with the image records recovered by the partial catalog recovery with the disk media IDs present in the DR domain. For this, run the command

```
nbcatsync -backupid <restored catalog backup ID>
```

DRAFT

# About site loss protection with catalog replication

This chapter includes the following topics:

- [About auto image replication](#)
- [About conditions for support of replicated NetBackup catalogs](#)
- [About catalog synchronization](#)
- [About NetBackup catalog replication](#)

## About auto image replication

The Auto Image Replication feature allows backups to be duplicated between the NetBackup domains and it automatically creates the catalog entries in the target domain as the backups are duplicated. Symantec recommends the use of Auto Image Replication instead of live catalog replication as a means of populating the NetBackup catalog at a disaster recovery site. Refer relevant section in NetBackup Administrator's Guide for more information on Auto Image Replication. The document discusses alternative methods of replicating the catalog data are in case the network environment does not lend itself to the use of Auto Image Replication.

## About conditions for support of replicated NetBackup catalogs

A NetBackup environment set up for replication is supported in the same way as any other NetBackup server. If the replicated catalog volume fails and is unrecoverable within a reasonable amount of time, NetBackup support

recommendations are the same as in the case of an unrecoverable disk failure of a non-replicated catalog. You should restore the catalog from the latest available catalog backup on the primary master server.

---

**Note:** Data can be lost in any data replication solution. To protect the NetBackup catalog, you must not solely rely on the replication technology due to the risk of failure of the replication technology. Data on the primary NetBackup server can get corrupted due to replication to the secondary hot standby NetBackup server. Therefore, you must frequently back up the NetBackup server catalogs.

---

---

**Warning:** Replication can adversely affect the application performance. Since additional time is required to commit changes to the NetBackup catalog, it may affect the overall backup times. Use replication at your own risk. Symantec shall have no liability for any replication failure based on your failure to properly install, configure, and monitor your replication solution.

---

The conditions of support for replication of NetBackup catalogs are as follows:

- The replication technology that is employed must maintain a consistent and write-ordered copy of the data at all times.
- The use of asynchronous replication technologies is allowed, if write-order fidelity can be maintained.
- The use of scheduled replication technologies such as hourly snapshots is not supported.
- The NetBackup master server and the EMM server must reside on the same virtual server that is controlled as a single entity.
- The primary and the secondary master servers must be of similar type, specification, operating system, and use the same virtual host name.
- The secondary master server must not have any other NetBackup function, neither in the same domain as the primary master server, nor in another domain. For example, you cannot use the secondary master server as a media server if it is not used as a master server. You also cannot use it as a master server for another NetBackup domain. Catalogs are replicated but cannot be merged.
- All master servers must be running NetBackup version 6.5.4 or higher.
- Configure both the clustered and the non-clustered environments to use a virtual hostname and IP address for the NetBackup master server that is separate from the physical host names and IP addresses of the servers. Separate virtual hostname and IP address let you control the active master server node

through DNS routing. It also prevents the primary and the secondary master servers from being active in the domain at the same time. For clustered environments this requirement is met automatically by the cluster configuration. For non-clustered environments the virtual hostname must be specified during installation.

- Ensure that the primary master server and the secondary master server use the same version of NetBackup and dependent component. Verify that the operating system, NetBackup binaries, EEBs, and configurations files that are not included in the paths are specified for replication.
- Replication between clustered and non-clustered master servers is not possible. Server pairs must be either clustered or non-clustered.
- The NetBackup catalog mount point must be the same at both the primary and the secondary sites.
- Only the catalog data is replicated between servers and must all be co-located on a single volume or volume set for replication. For clustered master servers the cluster common volume is replicated.  
 For non-clustered master servers, for details of the paths that must be linked to a volume set for replication, See [“About non-clustered NetBackup master server with catalog replication”](#) on page 58.
- Ensure that the virtual name or DNS alias does not resolve to both the primary and the secondary hosts at the same time.
- Catalog replication does not remove the requirement for catalog backup. Regularly back up the NetBackup catalog from the primary master server to protect against accidental image expiration or other inconsistencies that are introduced in the catalog on the primary site and replicated to the secondary site.
- If catalogs are replicated between NetBackup domains (rather than to a secondary server that can access the primary domain’s media servers) only the backups that are written to the tape and the replicated BasicDisk storage can be restored in the disaster recovery domain.
- Replication of the catalogs to a secondary master server lets you restore data during a short-term outage of the primary master server. In cross domain replication configurations, ensure that backups can be run after a failover. The catalogs should be able to be failed back to the primary server at a later date without data loss. Consider this support condition when you plan making backups at the DR site during a prolonged outage and then moving back to the primary site without losing information about the backups that are created at the DR site.

- Verify if NetBackup comes up using the replicated copy on the secondary site. This usage is not a requirement for support.
- Both the catalog and the backup images must be accessible at the secondary site.  
Users need to address the procedures that are related to availability of valid copies of the backup images. Users should also define procedures for enabling the NetBackup server to restore from the images at the secondary site. This document does not address these procedures.
- Users are responsible for installing, configuring, and monitoring their data replication solution. Users must ensure that the replication technology continuously maintains a consistent write-ordered copy of the NetBackup catalog volume.

## About catalog synchronization

Replication is a near instantaneous activity compared to the movement of tapes between sites. Replicated catalog data that is presented in the DR domain can be more current than the stock of tapes available in the DR domain which are dispatched from the production domain some time earlier. During restore operations, select only the backups that are created before the tapes were dispatched from the production domain for restore.

## About NetBackup catalog replication

To decide the NetBackup data protection strategy, you need to decide whether the DR site should be part of the same NetBackup domain or be a separate NetBackup domain.

NetBackup can be configured with catalog replication in the following ways:

- Multi-site single domain replication  
See [“About multi-site single domain replication”](#) on page 44.
- Multi-site cross domain replication  
See [“About multi-site cross domain replication”](#) on page 47.

## About multi-site single domain replication

Multi-site single domain is used where clients and media servers at both sites are under control of a common master server. Since both servers are part of the same domain, they see the same media servers and clients, and the NetBackup catalog is completely valid on the secondary master server.

In the multi-site single domain model, NetBackup catalogs are replicated between the sites. In the event of a problem at the primary site, the master server is failed over to a standby node on the secondary site. Backups are created on both sites (either by in-line copy or duplication depending on the configuration). Thus, the loss of a single site does not represent a true disaster, but loss of a number of application servers. Because the backup domain spans both sites, the loss of a single site results in reduction of the backup and restore capability, rather than destroying the backup environment. The multi-site single domain model uses a combination of master server clustering and storage replication. This combination allows the master server to be relocated easily and quickly to the secondary location.

The multi-site single domain model can be configured in following ways:

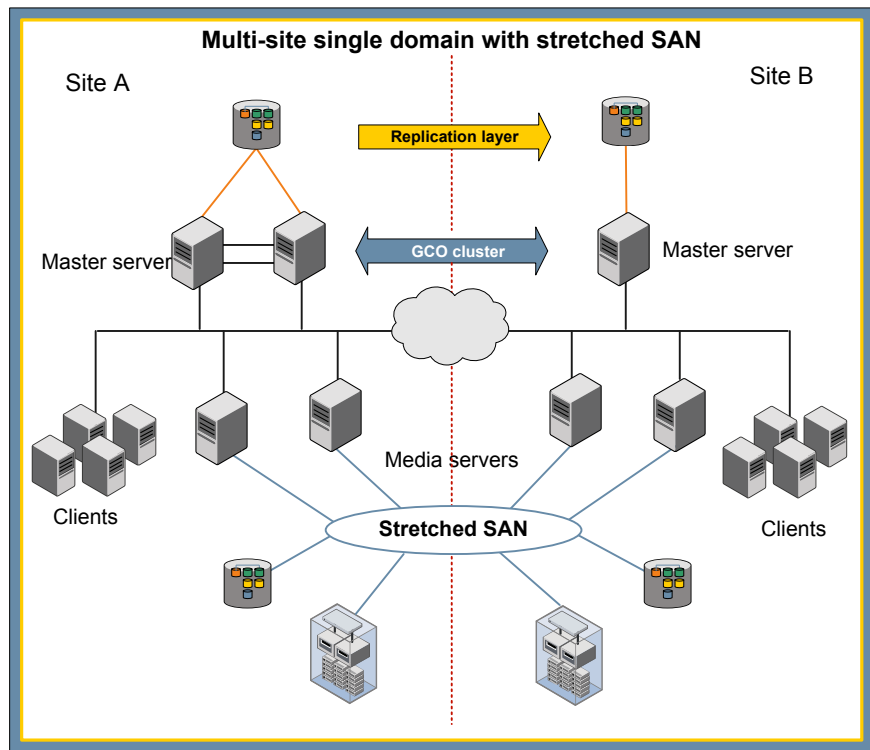
- Multi-site single domain with stretched SAN  
See “[About multi-site single domain with stretched SAN](#)” on page 45.
- Multi-site single domain with optimized duplication  
See “[About multi-site single domain with optimized duplication](#)” on page 46.

## About multi-site single domain with stretched SAN

To configure a multi-site single domain with stretched SAN, the media servers at each site must be configured with SAN access to backup devices at both sites. This access allows media servers to write and duplicate backups between the sites. This configuration works well for distances of up to 50 miles between sites, but becomes less effective as distance and latency increase.

[Figure 4-1](#) displays how a replicated global cluster is configured with multi-site single domain with stretched SAN.

Figure 4-1 Multi-site single domain with stretched SAN

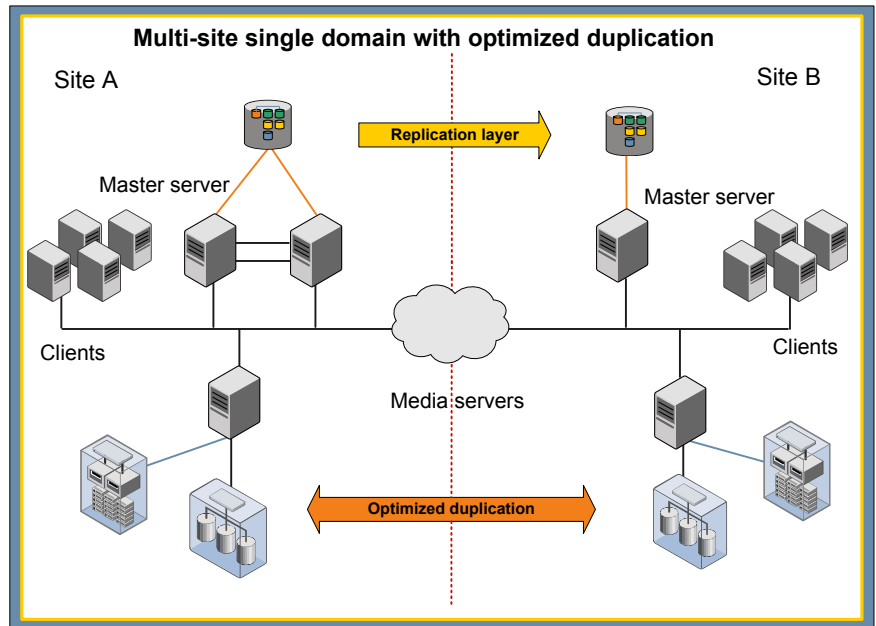


### About multi-site single domain with optimized duplication

To configure a multi-site single domain with optimized duplication, the stretched SAN must be replaced with a connection between the OpenStorage devices carrying out optimized duplication. In this configuration, the geographical separation can be greater because smaller data volumes are exchanged between sites. Using the hierarchical duplication capability in storage lifecycle policies, it is possible to create backups in the OpenStorage device at one site. You can then duplicate the backups to the OpenStorage device at the other site and finally duplicate the duplicated copy to tape for long-term storage.

Figure 4-2 displays how a replicated global cluster is configured with multi-site single domain with optimized duplication.

**Figure 4-2** Multi-site single domain with optimized duplication



## About multi-site cross domain replication

Multi-site cross domain replication is used when the DR site is a separate NetBackup domain than the production domain. The DR site has different media servers and devices.

Multi-site cross domain replication is only supported for tape and BasicDisk storage. AdvancedDisk types have specific media server or device configuration requirements that do not allow them to be accessed in disaster recovery domains.

## About multi-site cross domain and BasicDisk storage

You can replicate the images that are stored on non-staging BasicDisk storage between domains. The replication target must be mounted against the same mount point on a media server in the DR domain. Also, set the `FAILOVER_RESTORE_MEDIA_SERVER` parameter to ensure that the correct media server is selected. For example, you can replicate a BasicDisk storage unit using the mount point `/BD1` on media server `prdmed1` in the production domain to the DR domain. `/BD1` can be mounted on a media server `drmed1`, if the `bp.conf` file on the DR master server is edited to set `FAILOVER_RESTORE_MEDIA_SERVER = prdmed1 drmed1`. This setting is only possible for the BasicDisk storage units

that do not act as staging storage units and are not supported with staging storage units or other disk types.

### Planning a cross domain replication disaster recovery domain

To use the replicated catalog data on the secondary master server in the DR domain, ensure that the master server, media servers, network connections, and NetBackup software are functional.

Symantec recommends that you document the DR configuration steps, particularly if the DR domain is not normally configured. This documentation is particularly important if the domain is a facility provided by a specialist DR services company. When you are preparing a DR plan, the following points must be considered:

**Table 4-1** Planning a cross domain replication disaster recovery domain

Step	Description
Step 1	<p>Install the same NetBackup version on the master server, media servers, and clients in the DR domain that is used at the production domain.</p> <p><b>Note:</b> If the production domain has media servers with older versions of NetBackup, do not install the older version on the media servers in the DR domain. Use the same version for the master server and the media servers in the DR domain.</p> <p><b>Note:</b> If the full catalog replication method is used and the master server at the production domain is clustered, a clustered master server must also exist in the DR domain. The member nodes of the cluster do not need to be the same as those nodes at the production domain. If the partial catalog replication method is used then a clustered master server in the DR domain is not required.</p>
Step 2	<p>Test the network connectivity and authentication between the clients and servers using test backup policies. Disable the policies after testing.</p>
Step 3	<p>Tape drives and libraries must be connected to the media servers. The tape drives used in the DR domain must be read-compatible with the tapes from the production domain. They must be configured as the same media type in NetBackup.</p>
Step 4	<p>Set the <code>FAILOVER_RESTORE_MEDIA_SERVER</code> parameter to allow backups to be written to the media servers at the production domain so that backups can be restored using the media servers in the DR domain.</p>
Step 5	<p>If partial replication method is used, create a non-scratch Media Pool which is not used by any backup policy. Configure bar code rules to ensure that the backup tapes are automatically added to that pool.</p>

**Table 4-1** Planning a cross domain replication disaster recovery domain  
(continued)

Step	Description
Step 6	If different library types are used in the DR domain and at the production domain, ensure that the barcode masking operates in the same way. Remove the trailing characters wherever required. You can configure rules to manage this operation.
Step 7	<p>Ensure the following:</p> <ul style="list-style-type: none"><li>■ If the original backup tapes are used for DR purposes, they must be loaded in the tape libraries in the DR domain.</li><li>■ If backups are duplicated to secondary tapes for DR purposes, then load the off-site tapes in the tape libraries. Also the ALT_RESTORE_COPY_NUMBER file is created with the appropriate copy number in it.</li></ul> <p><b>Note:</b> Symantec recommends that the tapes are physically write-locked before they are placed in libraries in the DR domain. This locking reduces the risk of accidental overwriting of valid backups.</p>

## About full catalog replication

In full catalog replication, all parts of the catalog are replicated to the secondary master server. In full catalog replication, the tape information from the production domain, the media pool, and other assignments is retained. Backups can be run in the DR domain using the same policies and tapes that are used at the production domain. The replication can be reversed, which simplifies a transition back to the production domain. However, replicating the relational database components implies that the device configuration and the server configuration of the production domain is replicated to the DR domain. This configuration information cannot be used and the configuration in the DR domain must be discovered after recovery.

Full catalog replication is not recommended for cross domain replication.

## Recovering the catalog with full catalog replication

With full catalog replication, complete catalog backup is recovered to the DR master server. The media servers that do not exist in the DR environment should be deactivated to avoid unnecessary pooling. Since the device configuration at the DR site is likely to be different to the production site all device records are removed. Further, device discovery is run to update the EMM database.

This approach assumes that NetBackup is installed but not running on the secondary master server and the media servers in the DR domain. Also, the

secondary master server and the media servers are configured to communicate with each other.

Before restores can be started, carry out the following procedure to prepare for full catalog restore. You must document this procedure in your DR plan:

- 1 Ensure that replication between the primary and the secondary sites is stopped.

The replication is stopped if the primary master server is unavailable or if the replication link is disabled.

- 2 Mount the replicated volume to the appropriate mount point on the secondary master server.
- 3 Start the NetBackup Relational Database Manager, NetBackup PBX, and EMM services on the new master server.

■ On UNIX and Linux master servers run the following commands:

- `/usr/opensv/netbackup/bin/nbdbms_start_stop start`
- `/opt/VRTSspb/bin/pbx_exchange`
- `"/usr/opensv/netbackup/bin/nbemmm -maintenance`

■ On Windows master servers start the following Windows services:

- NetBackup Relational Database Manager
- Symantec Private Branch Exchange
- NetBackup Enterprise Media Manager

---

**Note:** The PBX process may already be running since it is not stopped and started by the NetBackup startup and shutdown commands.

---

- 4 Deactivate the media servers that are not part of the DR environment. Run the following command:

```
nbemmcmd -updatehost -machinename <Media Server> -machinestateop  
set_admin_pause -machinetype media -masterserver <Master Server>
```

- 5 If any media servers in the DR domain have the same names as media servers in the production domain, delete all tape devices from the EMM database. Run the following command:

```
nbemmcmd -deletealldevices -allrecords
```

---

**Note:** This step resolves possible device configuration conflicts on media servers. Skip this step, if the media servers in the DR domain have different names to those of the media servers in the production domain.

---

- 6 Restart NetBackup.
- 7 Optionally, you can deactivate all backup policies to prevent backups from starting automatically.
  - You can deactivate the backup policies **manually** using the **NetBackup Administration Console**.
  - Or run the `bppllist <policy> -set -inactive` CLI.
- 8 Register the media servers that form part of the DR environment in EMM by starting NetBackup on each media server.
- 9 Using the **Device Configuration Wizard**, create the new tape drive and library configuration.
- 10 Using the **NetBackup Administration Console**, verify if all the recovery media are set to non-robotic.
- 11 If some recovery media still need to be set to non-robotic, do the following:
  - Select the robotic media, right-click, and select **Move**.
  - Change the robot field to **Standalone**.
  - Click **OK** to save the changes.
- 12 Once all the recovery media are set to non-robotic, in the **Inventory all the tapelibraries** field ensure that the media are identified in the correct library.

You can now start restore and recovery operations of client data that is backed up at the production datacenter.

## Making the DR environment consistent with full catalog replication

In the event of a major incident at the production site, operate from the DR site for some time after the basic recovery is completed. The following additional tasks

may be optionally carried out once the DR environment is operational to make the DR environment consistent.

#### To make the DR environment consistent

- 1 Modify and enable the catalog backup policy and the other backup policies that are required in the DR domain.
- 2 Delete the policies that are no longer required.
- 3 Delete the storage units that are associated with the media servers that do not form part of the DR environment.

## About partial catalog replication

In partial catalog replication only the image database, policy, and the client configuration are replicated and the relational database components are not replicated. This allows the media servers and the devices to be preconfigured in the disaster recovery domain. You do not need to rediscover them in the event of a failover to the secondary master server.

As partial catalog replication does not replicate the relational database components of the NetBackup catalog, additional steps are required to be carried out following a failover to the disaster recovery master server before backups can be restored.

### Preparing an environment for partial catalog replication

The catalog image metadata, which is required to run restore operations, is stored in the relational database so a backup of the relational database must be taken at regular intervals and replicated along with the flat file information.

- 1 Change the configuration on the source (production) master server to ensure that the staging area for the relational database is located on the replicated storage. It can be done as follows:

- Create a suitable directory on the replicated storage.
- Use the following command to make this directory the staging area.

```
nbdb_admin -vxdbms_nb_staging <directory>
```

- 2 Backup the relational database to the staging area several times per day (ideally hourly) by running the following command in a scheduled script.

```
nbdb_backup -online <directory>-truncate_tlog
```

### Recovering the environment with partial catalog replication

In the event of a loss of the source master server (or during a disaster recovery test) follow these steps:

- 1 Ensure that replication between the primary and the secondary sites is stopped.  
  
Replication stops if the primary master server is unavailable or if the replication link is disabled.
- 2 Mount the replicated volume to the appropriate mount point on the secondary master server.
- 3 Use the command `nbdb_admin -vxdbms_nb_staging <directory>` on the target (disaster recovery) master server to point the staging area for the relational database to the location on the replicated storage.
- 4 Run the command `cat_export -all -staging` to export the metadata from the replicated relational database backup.
- 5 Run the command `cat_import -all` to import the exported metadata into the active relational database.
- 6 Start NetBackup on the secondary master server.
- 7 If the backup policies are replicated, deactivate all backup policies to prevent backups from starting automatically.
  - You can deactivate the backup policies manually using the **NetBackup Administration Console**.
  - Or run the command `bppllist <policy> -set -inactive`.
- 8 Ensure that the appropriate `FAILOVER_RESTORE_MEDIA_SERVER` settings are defined to direct restore operations through the media servers at the secondary site.

9 In order to restore backups from tapes the tapes must be added to the disaster recovery master server's catalog by placing them in a tape library and running an inventory of the library. To prevent the tapes from being accidentally overwritten the disaster recovery master server should have a bar code rule that adds the tapes to a volume pool that is not the global scratch pool and is not used by any backup policies. Ideally the tapes should also be physically write locked.

10 For disk based backups, the storage servers and disk pools must be added to the disaster recovery master server by running the disk storage server wizard.

Once the disk storage is present, run the following command to reconcile the disk media IDs:

```
nbcatsync -backupid <catalog backup ID> -prune_catalog
```

The value <catalog backup ID> is the backup ID of the most recent catalog backup and can be found in the catalog backup's disaster recovery file. Once the tapes have been added and the disk media IDs have been reconciled it is possible to start restore operations

## Making the disaster recovery environment consistent with partial catalog replication

In the event of a major incident at the production site, operate from the disaster recovery site for some time after the recovery is completed. The following additional tasks may be optionally carried out once the disaster recovery environment is operational to make the disaster recovery environment consistent.

### To make the disaster recovery environment consistent with partial catalog replication

- 1 Modify and enable the catalog backup policy and any other backup policies that are required in the disaster recovery domain.
- 2 Delete the policies that are no longer required.

## Considerations for managing tapes with partial catalog replication

The tapes from the production domain are not assigned in the disaster recovery domain. The tapes must be manually added to the database and placed in a pool where they cannot get accidentally overwritten. This can also be done using a combination of barcode rules and the robot inventory command.

As the tapes are not assigned on the disaster recovery master server they will not be released to the global scratch pool when backups expire and therefore these tapes must be manually recycled.

---

**Caution:** Care must be taken to ensure that the tapes are manually moved to the global scratch pool only when they do not have valid backups on them.

---

The simplest way of checking this is to create two lists by running the commands `bpimagelist -d "01/01/1970 00:00:00" -media -l` and `vmquery -pn <private pool name> -b` and then comparing the lists. Tapes found in the second list but not found in the first list have no valid images on them and can be moved to the scratch pool by running the command `vmchange -p <scratch pool number> -m <media id>`.

DRAFT

DRAFT

# Deploying NetBackup master servers with full catalog replication

This chapter includes the following topics:

- [About replication considerations](#)
- [About non-clustered NetBackup master server with catalog replication](#)
- [About globally clustered NetBackup master servers with catalog replication](#)

## About replication considerations

To deploy NetBackup with catalog replication, you must consider the following factors for planning the actual deployment.

**Table 5-1** Replication considerations

Considerations	Description
Master server considerations	<p>Symantec does not recommend operating a master server as a combined master and media server. If the storage devices available at the different sites are not compatible, it can lead to problems with storage unit definitions and backup failures.</p> <p>Catalog replication is not a substitute for catalog backup and the catalog must be backed up on a regular basis.</p>

**Table 5-1** Replication considerations (*continued*)

Considerations	Description
Networking considerations	<p>In a multi-site single domain configuration, the master server controls the media servers on both the sites. The metadata must pass between the sites. This metadata traffic is sent over a standard I/P link between the sites. The same link can be used as the heartbeat link for the global cluster control. Symantec recommends that a link of at least 10 Mb/sec and ideally 100 Mb/sec must be provided between the sites to handle this traffic.</p> <p>If host-based replication is used, additional I/P bandwidth is required for the replication layer. The additional bandwidth must also be factored in.</p>
DNS considerations	<p>If the master server nodes at the secondary site are on a different subnet from the master server nodes at the primary site, a DNS change is required as part of the failover process. You can initiate the DNS change automatically by using the cluster failover process. You can also initiate the process manually. The backup system does not function correctly until the change is fully propagated, which can affect the recovery time in a site failover.</p> <p><b>Note:</b> To propagate the DNS change automatically by the cluster service group, the DNS resource must come on-line after NetBackup starts.</p>

## About non-clustered NetBackup master server with catalog replication

The following topic provides guidelines for installing, configuring, and operating non-clustered NetBackup master server cluster with catalog replication:

### Installing and configuring non-clustered NetBackup master server with catalog replication

The installation and configuration of non-clustered NetBackup master servers with catalog replication progresses through multiple stages which are described in [Table 5-2](#)

**Note:** VxSS or NBAC is not supported with catalog replication for non-clustered master servers. For more information about NBAC, refer to the *NetBackup Security and Encryption Guide*.

**Table 5-2** Installing and configuring non-clustered NetBackup master server cluster with catalog replication

Step	Description
Step 1	Make sure that you meet all the conditions of support before proceeding with the actual installation.  See <a href="#">“About conditions for support of replicated NetBackup catalogs”</a> on page 41.
Step 2	Install and configure the non-clustered NetBackup master server on the primary site.  See <a href="#">“Installing and configuring the primary NetBackup master server”</a> on page 59.
Step 3	Install and configure the non-clustered NetBackup master server on the secondary site.  See <a href="#">“Installing and configuring secondary NetBackup master server”</a> on page 61.

## Installing and configuring the primary NetBackup master server

The primary master server is the server that normally functions as the master server. Thus, it must be installed first.

The following procedure provides guidelines for installing and configuring the primary non-clustered master server with catalog replication.

### To install and configure the primary non-clustered master server with catalog replication

- 1 You must use a DNS alias name for the master server. The DNS alias name ensures a smooth failover to the secondary master server. Before you start the installation, define this alias name in DNS and map it to the primary master server. Configure all the media servers and clients in the NetBackup domain to use this alias name for the master server and EMM server.
- 2 Install the NetBackup master server on the primary master server node. Specify the alias name for the master server and EMM server.
- 3 Shut down NetBackup after the installation is complete.
- 4 To ensure that NetBackup starts correctly when switching to the secondary master server, modify the `server.conf` and `vxdbsms.conf` files. Also, change the string `NB_<hostname>` to `NB_<alias name>` in each file.

On Windows master server, these files are located at:

■ `<install path>\VERITAS\netbackupdb\conf\server.conf`

- `<install path>\VERITAS\netbackupdb\data\vxdbms.conf`

On UNIX and Linux master server, these files are located at:

- `/usr/openv/var/global/server.conf`
- `/usr/openv/db/data/vxdbms.conf`

**5 Move the catalog components to the volume that is replicated to the secondary master server.**

For Windows installations, map the following paths to a common volume. Use the folder mount points or junction points on Windows 2003 and symbolic links on Windows 2008:

- `<install path>\VERITAS\netbackup\db`
- `<install path>\VERITAS\netbackupdb\data`
- `<install path>\VERITAS\netbackup\vault\sessions`
- `<install path>\VERITAS\volmgr\misc`
- `<install path>\VERITAS\netbackup\var`
- `<install path>\VERITAS\kms`

For UNIX and Linux installations, soft link the following paths to locations on a common volume:

- `/usr/openv/netbackup/db`
- `/usr/openv/db/data`
- `/usr/openv/netbackup/vault/sessions`
- `/usr/openv/volmgr/database`
- `/usr/openv/var`
- `/usr/openv/kms`

---

**Note:** You can use the `nbdb_move` command to relocate the EMM database rather than linking `/usr/openv/db/data` for UNIX and Linux or `<install path>\VERITAS\netbackupdb\data` for Windows. However, all other paths must be linked.

---

- 6 Configure NetBackup so that it can be manually started and stopped on the primary master server. By default, NetBackup is started automatically when the master server is booted.

To prevent this automatic start, make the changes as per Steps 7, 8 and , 9 after the initial installation and after applying patches or upgrades.

- 7 Make the following changes on the NetBackup master server.

- On UNIX and Linux master server, remove the links to `/etc/init.d/netbackup` created during the installation to enable automatic startup.  
Refer to the *Symantec NetBackup Installation Guide* for details of links for each operating system.

- On Windows master server, go to the **Services Manager** and set the **Startup type** for all the NetBackup services to **Manual**.

- 8 Start NetBackup on the primary master server to confirm that it comes up correctly and then shut it down again.

At this stage, you may run the **NetBackup Configuration Wizard** to configure media servers and storage devices.

- 9 To start and stop NetBackup manually run the following commands. It is recommended to document these commands in the failover procedure.

On UNIX and Linux master server:

- To start NetBackup, run the following command:  
`/etc/init.d/netbackup start`
- To stop NetBackup, run the following command  
`/etc/init.d/netbackup stop`

On Windows master server:

- To start NetBackup, run the following command:  
`<install path>\VERITAS\NetBackup\bin\bpup`
- To stop NetBackup, run the following command:  
`<install path>\VERITAS\NetBackup\bin\bpdown`

## Installing and configuring secondary NetBackup master server

The following procedure provides guidelines for installing and configuring the secondary non-clustered master server with catalog replication.

**To install and configure the secondary non-clustered master server with catalog replication**

- 1 Stop NetBackup on the primary master server.
- 2 Map the DNS alias name to the secondary master server.
- 3 Install the NetBackup master server on the secondary master server node, specifying the alias name for the master server and EMM server. During installation, apply the same list of servers on the secondary master server.
- 4 After the installation is complete, shut down NetBackup.
- 5 To ensure NetBackup starts correctly when switching to the secondary master server, modify the `server.conf` and `vxdbms.conf`. Also, change the string `NB_<hostname>` to `NB_<alias name>` in each file.

On Windows master server these files are located at:

- `<install path>\VERITAS\netbackupdb\conf\server.conf`
- `<install path>\VERITAS\netbackupdb\data\vxdbms.conf`

On UNIX or Linux master server these files are located at:

- `/usr/opensv/var/global/server.conf`
- `/usr/opensv/db/data/vxdbms.conf`

- 6 Create a small disk volume (100 MB) and mount it to the same mount point used for the replicated volume on the primary server. This volume is known as the upgrade volume and is to be presented every time NetBackup is upgraded on the secondary master server.

---

**Note:** During a failover operation the replicated volume is mounted on the secondary master server and not the upgrade volume.

---

- 7 Move the catalog components to the upgrade volume.

For Windows installations map the following paths to a common volume using folder mount points or junction points on Windows 2003 and symbolic links on Windows 2008:

- `<install path>\VERITAS\netbackup\db`
- `<install path>\VERITAS\netbackupdb\data`
- `<install path>\VERITAS\netbackup\vault\sessions`
- `<install path>\VERITAS\volmgr\misc`
- `<install path>\VERITAS\netbackup\var`

- <install path>\VERITAS\kms

For UNIX and Linux installations, soft-link the following paths to locations on a common volume:

- /usr/openv/netbackup/db
- /usr/openv/db/data
- /usr/openv/netbackup/vault/sessions
- /usr/openv/volmgr/database
- /usr/openv/var
- /usr/openv/kms

- 8 Configure NetBackup so that it can be manually started and stopped on the primary master server. By default, NetBackup is started automatically when the master server is booted.

To prevent this automatic start, make the changes as per Steps 7, 8 and , 9 after the initial installation and after applying patches or upgrades.

- 9 Make the following changes on the NetBackup master server.

- On UNIX and Linux master server, remove the links to /etc/init.d/netbackup created during the installation to enable automatic startup.  
Refer to the *Symantec NetBackup Installation Guide* for details of links for each operating system.

- On Windows master server, go to the **Services Manager** and set the **Startup type** for all the NetBackup services to **Manual**.

- 10 Start NetBackup on the primary master server. Confirm that NetBackup comes up and then shut it down again.

At this stage, you may run the **NetBackup Configuration Wizard** to configure media servers and storage devices.

- 11 To start and stop NetBackup manually the run the following commands. It is recommended to document these commands in the failover procedure.

On UNIX and Linux master server:

- To start NetBackup, run the following command:  
/etc/init.d/netbackup start command
- To stop NetBackup, run the following command:  
/etc/init.d/netbackup stop

On Windows master server:

- To start NetBackup, run the following command:  
`<install path>\VERITAS\NetBackup\bin\bpup`
- To stop NetBackup, run the following command:  
`<install path>\VERITAS\NetBackup\bin\bpdown`

12 After NetBackup shuts down, dismount the upgrade volume. Then reset the DNS alias name to the primary master server. Then restart NetBackup on the primary master server.

### Upgrading NetBackup master server in a non-clustered replicated configuration

To upgrade NetBackup in a replicated configuration, the primary and the secondary master servers must run the same version of NetBackup. If the secondary master server is not running the same version as the primary master server, NetBackup may not start in a failover situation. The primary master server is upgraded in the usual way. However it is recommended that replication of the catalog is suspended while the upgrade is carried out. Refer to the *Symantec NetBackup Installation Guide* for detailed instructions.

Guidelines for upgrading the secondary NetBackup master server in a replicated configuration are described in the following table.

Table 5-3 Upgrading NetBackup master server in a non-clustered replicated configuration

Step	Description
Step 1	Present the upgrade volume to the appropriate mount point.
Step 2	Reassign the DNS alias. You can also create a temporary local alias for the master server.
Step 3	Start NetBackup.
Step 4	Install the latest version of NetBackup.
Step 5	Stop NetBackup.
Step 6	Reassign the DNS alias. Or remove the local alias.
Step 7	Dismount the upgrade volume.

**Note:** During major NetBackup upgrades such as, from 6.5.x to 7.0 the `server.conf` and `vxdbsms.conf` files may be replaced. After upgrading always check these files. Ensure that they contain the correct name for the master server and modify them if necessary.

## About globally clustered NetBackup master servers with catalog replication

This section provides guidelines for installing, configuring, and operating a globally clustered NetBackup master server with catalog replication.

### Installing and configuring a globally clustered NetBackup master server with catalog replication

The installation and configuration of clustered NetBackup master servers with catalog replication progresses through multiple stages which are described in [Table 5-4](#).

**Table 5-4** Installing and configuring clustered NetBackup master server cluster with catalog replication

Stage	Description	Action
Stage 1	Installation prerequisites	<p>Make sure that you meet all the conditions of support before proceeding with the actual installation.</p> <p>See <a href="#">“About conditions for support of replicated NetBackup catalogs”</a> on page 41.</p> <p>See <a href="#">“About replication considerations”</a> on page 57.</p>
Stage 2	Installing and configuring the clustered NetBackup master on the primary site	<p>Install and configure the NetBackup master server cluster on the primary site.</p> <p>See <a href="#">“Installing and configuring primary NetBackup master server cluster”</a> on page 66.</p>
Stage 3	Installing and configuring the clustered NetBackup master on the secondary node	<p>Install and configure the NetBackup master server cluster on the secondary site.</p> <p>See <a href="#">“Installing and configuring secondary NetBackup master server cluster”</a> on page 67.</p>

### About clustering considerations

The NetBackup master server nodes on both the sites must be configured as clustered master servers, although they can be single node clusters on each site. NetBackup allows a member node to be upgraded only while the node is off line without affecting the database structure. Therefore, if master servers are not clustered, there is a risk of data loss and corruption during an upgrade operation.

The NetBackup master server is an active or a passive application. It can only run on one node of the cluster at any one time. In a replicated environment, the cluster members at both sites effectively form a single cluster. You can create a global cluster with two to four nodes, depending on the level of resilience required.

Single node cluster at both sites	This configuration requires two nodes—one node at each site. This configuration is the most efficient in terms of the servers involved. The disadvantage of this configuration is that even a local problem with the primary master server requires a site failover operation.
Dual node on primary site and single node on secondary site	<p>This configuration requires three nodes—two at the primary site and one at the secondary site. During normal operations, since there is single node at each site, site failover is not required to address issues with the primary master server. Instead a local failover can be used. However, there is no protection for the node on the secondary site.</p> <p>As a general best practice, this configuration is recommended.</p>
Dual node on both sites	This configuration requires four nodes, three of which are always idle. This configuration allows local failover capability at sites. With this configuration, if a local server problem is encountered, there is no need to failover.

### Installing and configuring primary NetBackup master server cluster

To install the NetBackup master server cluster follow the instructions that are described in the *Symantec NetBackup Clustered Master Server Administrator's Guide*. Guidelines to install the primary NetBackup master server cluster with catalog replication are described in the following table.

**Table 5-5** Guidelines to installing the primary NetBackup master server cluster with catalog replication

Step	Description
Step 1	When you install NetBackup master server cluster on primary node, specify: <ul style="list-style-type: none"> <li>■ Replicated storage as the mount point for the cluster common storage.</li> <li>■ All servers that are part of the domain.</li> <li>■ Servers that form the secondary site cluster.</li> </ul>
Step 2	After the NetBackup cluster group is created, reconfigure the storage resources to include the replication control components.
Step 3	For some replications layers, especially for Veritas Volume Replicator (VVR), the replication agent must be in a separate service group. You must link the agent with the NetBackup application service group.
Step 4	If the replication technology offers a bandwidth planning and analysis tool, configure the NetBackup master server cluster to run on the primary master server cluster for some time. Perform the configuration before you implement the replication layer and stretch the cluster to the secondary site.

## Installing and configuring secondary NetBackup master server cluster

To install and configure the secondary NetBackup master server cluster, it is not necessary to have the same number of nodes as the primary NetBackup master server cluster. The secondary NetBackup master server cluster can be a single node cluster but should be clustered. Guidelines to install the secondary NetBackup master server cluster with catalog replication are described in the following table.

**Table 5-6** Guidelines to installing the secondary NetBackup master server cluster with catalog replication

Step	Description
Step 1	Install the secondary NetBackup master server cluster using the same virtual name and cluster common mount point as the primary NetBackup master server cluster.
Step 2	During installation, specify all the NetBackup servers, the master server cluster nodes and the media servers at the primary site.

**Table 5-6** Guidelines to installing the secondary NetBackup master server cluster with catalog replication (*continued*)

Step	Description
Step 3	Initially, to configure the secondary NetBackup master server cluster, use a temporary volume as the cluster common volume instead of replication target of the primary NetBackup master server cluster. A temporary volume helps to prevent data loss and access conflicts. The catalog setup that is done during Windows installation and cluster configuration on UNIX and Linux platforms creates and seeds the EMM database on the cluster common storage.
Step 4	<p>Create a small disk volume (100MB) and mount it to the same mount point used for replicated volume on the primary NetBackup master server cluster. After the cluster configuration is complete, take the cluster offline and map it with the common replication target. Remove its mapping from the temporary storage.</p> <p>This volume is required for Windows environments each time the secondary NetBackup master server cluster is upgraded. Thus, this volume is known as upgrade volume.</p> <p>For UNIX and Linux environments, this volume is not required after the initial installation. You can remove the upgrade volume after the installation.</p>
Step 5	Do not configure the replication agent on the secondary site to automatically reverse the direction of replication. Do not reverse the replication until the primary site is operational again.

## Populating the server tables in the EMM

The server tables must be correctly populated in the EMM. Server tables can be populated automatically by failing over the master server cluster to each node in turn.

Guidelines to populating the server tables in the EMM are described in the following table.

**Table 5-7** Guidelines to populating the server tables in the EMM

Step	Description
Step 1	After the cluster setup on the secondary site is complete, take NetBackup offline on the primary site.

**Table 5-7** Guidelines to populating the server tables in the EMM (*continued*)

Step	Description
Step 2	Reverse the replication direction and bring NetBackup on-line on the secondary site.  Bringing the NetBackup master server cluster on-line on a particular node of the cluster automatically adds that node to the list of known servers in the EMM server tables. In case of multiple nodes at the secondary site, the master server must be failed over to each node.
Step 3	After all the member nodes are added, take NetBackup offline from the secondary site.
Step 4	Reverse the replication direction and bring NetBackup online on the primary site.

## Upgrading NetBackup in clustered replication configuration

**Note:** The master server clusters at the primary and the secondary site must run the same version of NetBackup. If there are different versions, NetBackup may not start in a failover situation.

### Upgrading NetBackup in clustered replication configuration on UNIX and Linux

The procedure to upgrade NetBackup in the primary and the secondary master server cluster on UNIX and Linux platforms, is outlined in Symantec NetBackup Installation Guide. Also refer Symantec NetBackup Clustered Master Server Administrator's Guide.

**Note:** On UNIX and Linux platforms, all nodes on the secondary master server cluster are considered to be passive during the upgrade.

**To upgrade a UNIX or Linux global cluster, use the following steps:**

- 1 Ensure that global cluster failover is disabled during the upgrade. Suspend replication between the primary and the secondary sites.
- 2 Upgrade NetBackup on the primary site master server cluster first, following the standard upgrade procedures. Refer [NetBackup Upgrade Guide](#).

- 3 Ensure that the phase 1 and phase 2 image metadata migration processes have completed successfully before proceeding to upgrade the secondary site master server cluster.
- 4 After all the secondary site cluster nodes are upgraded, edit the **bp.conf** file on each node of the secondary site cluster and set the value of **LIST\_FS\_IMAGE\_HEADERS** to **NO**.
- 5 Restart the replication between the primary and the secondary sites and enable the global cluster failover.

## Upgrading NetBackup in clustered replication configuration on Windows

To upgrade the master server cluster on Windows, at least one node of the cluster must be active at the time of the upgrade even if there are no changes to the database schemas.

**To upgrade Windows primary master server cluster, use the following steps:**

- 1 Ensure that global cluster fail over is disabled for the duration of the upgrade.
- 2 Suspend replication between primary and secondary sites.
- 3 Upgrade NetBackup on the primary site master server cluster first, following the standard upgrade procedures. Refer [NetBackup Upgrade Guide](#).
- 4 Ensure that the phase 1 and phase 2 image metadata migration processes have completed successfully before proceeding to upgrade the secondary cluster.

**To upgrade Windows secondary master server cluster, use the following steps:**

- 1 Isolate the secondary master server cluster from the rest of the network.
- 2 Present the upgrade volume to the appropriate mount point. If no upgrade volume is available, see the note following this procedure.
- 3 Upgrade NetBackup using the standard upgrade procedures mentioned in the NetBackup Installation Guide.
- 4 Stop NetBackup.
- 5 On all nodes of the secondary cluster, ensure that the registry entry **HKEY\_LOCAL\_MACHINE** at **LIST\_FS\_IMAGE\_HEADERS** is set to **NO**.
- 6 Dismount the upgrade volume.
- 7 Reconnect the secondary master server cluster to the rest of the network.

**Note:** If no upgrade volume is available, consider using the following approach during the upgrade:

- Before upgrading the primary cluster, stop the replication between the primary and the secondary clusters so that the target volume can be mounted to the secondary cluster. (Refer the replication documentation).
- At step 2 of the secondary cluster upgrade procedure, present the replication target volume to the secondary cluster in place of the upgrade volume.
- After the upgrade of the secondary cluster is complete (step 7), re-establish the replication link and synchronize the target volume with the source volume. Take appropriate care while re-establishing the replication link to ensure that the synchronization occurs in the correct direction (from the primary site to the secondary site).

## Failing over to secondary master server cluster

You need to failover to the secondary master server cluster when:

- All the nodes of the primary master server cluster fail.
- Access to the primary site is denied.

The exact failover procedure can vary for different replication technologies.

[Table 5-8](#) lists the high-level steps to failover to the secondary master server cluster.

**Table 5-8** Failing over to secondary master server cluster

Step	Description
Step 1	Stop NetBackup on the primary master server cluster.
Step 2	Stop or reverse the replication of the catalog volume.
Step 3	If necessary, update the DNS with the new virtual IP address for the master server.
Step 4	Start NetBackup on the secondary master server cluster.

**Note:** Step 1 and 2 occur automatically if the primary site fails.

In globally clustered environment the process of failing over to secondary master server can be automated. To automate the process, multiple heartbeat connections must exist between the clusters. Failure of the heartbeat network can cause the

secondary master server cluster to come online while the primary master server cluster is still operational.

## Testing the NetBackup master server cluster in clustered replication environment

It is recommended to test the ability of the secondary master server cluster to come online without going through a full failover operation. In case of a full failover, the exact procedure can vary for different replication technologies.

Table 5-9 lists the high-level steps that must be followed for testing.

**Table 5-9** Testing the NetBackup master server cluster in clustered replication environment

Step	Description
Step 1	Suspend replication between the primary and the secondary master server clusters.
Step 2	Isolate the secondary master server cluster from the network.
Step 3	Start NetBackup on the secondary master server cluster.
Step 4	Perform the required validation checks. Ensure that NetBackup is running on the secondary master server cluster.
Step 5	Stop NetBackup on the secondary master server cluster.
Step 6	Re-establish network connections on the secondary master server cluster.
Step 7	Restart the replication.

# Using NetBackup to perform backups and restores in a cluster

This chapter includes the following topics:

- [About backups and restores with NetBackup in a cluster](#)
- [About supported NetBackup database agents and options in a cluster](#)

## About backups and restores with NetBackup in a cluster

This topic provides links to instructions on how to perform user-directed backups and restore data in a cluster. Additionally, specific instructions for performing backups and restores can be found in other NetBackup guides. See the *NetBackup Backup, Archive, and Restore Getting Started Guide* and the administrator guides for NetBackup's database agents and options.

The backup and restore process is the same whether you are in a cluster or a non-cluster environment. See the *Symantec NetBackup Troubleshooting Guide* for further information on backup and archive processes and on restore processes.

## Performing user-directed backups with NetBackup in a cluster

When you perform user-directed backups in a cluster, you can use the node name or the virtual name of the client to perform the backup. If you choose the virtual name, the backup can be restored from any of the cluster nodes. You can also configure automatic backups.

**To perform a user-directed backup on a Windows client**

- 1 Open the **Backup, Archive, and Restore** console.
- 2 On the **File** menu, click **Specify NetBackup Machines**.
- 3 From the **Source client** list, select (or add) the wanted node or virtual name.

**To perform a user-directed backup on a UNIX/Linux client**

- 1 Open the **Backup, Archive, and Restore** console.
- 2 In the **Login** dialog box, enter the name of the client, either the node or the virtual client name.

You must log on to the wanted node or virtual client. You cannot specify a client other than the local client in the Java interface.

## About restoring data in a cluster

For all file restore operations, use the procedures on how to perform restores in the *NetBackup Backup, Archive, and Restore Getting Started Guide*. When you restore files to the shared disk drives, restore those files to the virtual server name.

When you restore individual database files, restore those files to the virtual server name that corresponds to the client where the database application is installed.

---

**Note:** Since a computer can have multiple virtual names in a cluster environment, files can be backed up in the context of more than one client name. If you carefully plan your backup policies, you can avoid this problem. However, it may be necessary to browse more than one client name to locate a backup image. And you may need to perform more than one restore to restore all of the files that you need.

---

The Backup, Archive, and Restore console operates in the context of that client's name. You must perform a redirected restore to restore the files on the shared disk that were backed up with the virtual server name. NetBackup allows a redirected restore operation only if the necessary configuration is performed on the NetBackup master server. See the information on how to allow redirected restores in the *Symantec NetBackup Administrator's Guide, Volume I*.

There may be other situations that require the appropriate `altnames` directory entries to be created on the master server. While NetBackup tries to restore files from the client, the operation may fail with this error message:

```
131 client is not validated to use this server
```

If you see this message, you must set up the `altnames` directory to allow the operation to succeed. For example, the required network interface parameter may be set to a valid network name for the client. But this name may not match the NetBackup **Client name** parameter for that client. This situation often happens for NetBackup clients in a cluster. Alternatively, you can perform a server-directed restore and avoid the need to set up the `altnames` directory.

See [“Example: Performing a user-directed restore in a NetBackup cluster”](#) on page 75.

## Example: Performing a user-directed restore in a NetBackup cluster

For example, assume the cluster virtual server name is `TOE` and the cluster node names are `TIC` and `TAC`. Files on the shared disk must be backed up by a NetBackup policy that includes `TOE` in the client list.

To perform a server-directed restore of files on the shared disk, set both the source client and the destination client to `TOE`. The server-directed restore does not have to know which node is in control of the shared disk at the time of the restore.

### To perform a user-directed restore of files in a NetBackup cluster

- 1 Create the following files on the master server.

For a UNIX or Linux server:

```
/usr/opensv/netbackup/db/altnames/tic
/usr/opensv/netbackup/db/altnames/tac
```

For a Windows server:

```
shared_drive_install_path\NetBackup\db\altnames\tic
shared_drive_install_path\NetBackup\db\altnames\tac
```

- 2 In both files, add the virtual server name `TOE` on one line in the file.
- 3 Determine which node (`TIC` or `TAC`) has control of the shared disk.
- 4 Start the Backup, Archive, and Restore interface on that node and select the virtual server name (`TOE`) as the source client and the server.
  - On Windows computers, in the **File** menu, click **Specify NetBackup Machines**.
  - On UNIX or Linux computers, in the **Actions** menu, click **NetBackup Machines**.
- 5 Browse the backed-up files by using the virtual server name (`TOE`) from the shared disk and restore them as needed.

## About supported NetBackup database agents and options in a cluster

Only certain database agents and NetBackup options are supported in a clustered environment.

For information on how to install and configure database agents and options in a cluster, refer to the administrator's guide for that agent or option.

Backing up database files in a cluster Database applications are installed on a cluster as virtual servers. To protect the data for these virtual servers, install the appropriate NetBackup database agent on each node of the cluster. With NetBackup for Windows, database agents are installed along with NetBackup server and the NetBackup client.. Also create a backup policy for that database agent. When you configure a policy for the application or database in the cluster, always use the virtual server name of the application or database as the client name in the policy. For complete installation and configuration instructions for a particular database agent, see the NetBackup documentation for that agent.

User backups User backups that are run on individual nodes of the cluster generally run as a backup of the node, not the NetBackup virtual server. You may find it easier to use scheduled backups rather than user backups to protect the data in the cluster.

NetBackup client in a cluster You may choose to install only the NetBackup client in a cluster. In this configuration, you can back up the data from the cluster across the network to a separate NetBackup server. In this situation the NetBackup-specific configuration tasks for tape devices, media, and so on, are separate from the setup and maintenance of the cluster itself. However, the NetBackup client itself cannot fail over.

To install the NetBackup client on an MSCS, VCS, SunCluster, Service Guard cluster, or HACMP cluster

The NetBackup client is installed on a cluster as it is in a non-clustered environment. Refer to the *Symantec NetBackup Installation Guide* for information on how to install the NetBackup client. On Windows systems, you may have problems with name resolution when you try to back up data on the cluster. (This data can be either local data or shared data.) Consider setting the **Required Network Interface** parameter for each client to the fully qualified name of the node where the NetBackup client is installed.

## About backing up database files in a cluster

Database applications are installed on a cluster as virtual servers. To protect the data for these virtual servers, install the appropriate NetBackup database agent on each node of the cluster. With NetBackup for Windows, database agents are installed along with NetBackup server and the NetBackup client.. Also create a backup policy for that database agent. When you configure a policy for the application or database in the cluster, always use the virtual server name of the application or database as the client name in the policy. For complete installation and configuration instructions for a particular database agent, see the NetBackup documentation for that agent.

## About user backups

User backups that are run on individual nodes of the cluster generally run as a backup of the node, not the NetBackup virtual server. You may find easier to use scheduled backups rather than user backups to protect the data in the cluster.

## About NetBackup client in a cluster

You may choose to install only the NetBackup client in a cluster. In this configuration, you can back up the data from the cluster across the network to a separate NetBackup server. In this situation the NetBackup-specific configuration tasks for tape devices, media, and so on, are separate from the setup and

maintenance of the cluster itself. However, the NetBackup client itself cannot fail over.

To install the NetBackup client on an MSCS, VCS, SunCluster , Service Guard cluster, or HACMP cluster

The NetBackup client is installed on a cluster as it is in a non-clustered environment. Refer to the *Symantec NetBackup Installation Guide* for information on how to install the NetBackup client. On Windows systems, you may have problems with name resolution when you try to back up data on the cluster. (This data can be either local data or shared data.) Consider setting the **Required Network Interface** parameter for each client to the fully qualified name of the node where the NetBackup client is installed.

DRAFT

# Index

## A

adding disk storage unit 28

## B

backing up database files 77

backups

    user-directed 73

## C

catalog protection

    catalog replication 23

*See also* catalog replication

    online catalog backup 23

catalog recovery 31

    full catalog recovery 32

    partial catalog recovery 35

catalog replication

    catalog synchronization 44

    conditions for support 41

    considerations. *See* replication considerations

    full catalog replication 49

    multi-site cross domain replication 47

    multi-site single domain 44

    partial catalog replication 52

catalog synchronization 44

clustered master server

    installing primary master server 66

    installing secondary master server 67

    server table 68

    testing with catalog replication 72

    upgrade volume 67

    upgrading NetBackup 69

## D

database agents 76

dedicated media server protection

    storage unit group 18

disc recovery 37

## F

full catalog recovery 32

    restoring full catalog 33

full catalog replication 49

    restoring full catalog 49

full catalog restore 49

## L

LAN client protection 21

## M

master server protection

    clustering 17

media availability protection

    global scratch pool 16

    media sharing 16

media server

    non-failover cluster configuration 24

    restore backup 20

multi-site cross domain replication

    BasicDisk storage 47

multi-site single domain replication 44

    optimized duplication 46

    stretched SAN 45

## N

NetBackup client in cluster 77

network link protection

    redundant network teaming 15

non-clustered master server

    installing primary master server 59

    installing secondary master server 61

    upgrade volume 61

    upgrading NetBackup 64

non-dedicated media server protection

    storage unit group 19

non-failover cluster configuration

    media server 24

    SAN media server 24

**O**

options 76

**P**

partial catalog recovery 35  
     restoring partial catalog 36  
 partial catalog replication 52  
 point of failure  
     component 13  
     dedicated media server 18  
     LAN client 21  
     master server 17  
     media availability 16  
     media server 18  
     network link 15  
     non-dedicated media server 19  
     SAN client 22  
     SAN media server 20  
     site 22  
     storage device 16  
     storage device connection 15  
         robotic control connection 15  
         SAN connection 15  
 prerequisites standalone media server 25  
 protection method  
     dedicated media server  
         storage unit group 18  
     LAN client 21  
     master server  
         clustering 17  
     media availability  
         global scratch pool 16  
         media sharing 16  
     network link  
         redundant network teaming 15  
     non-dedicated media server  
         storage unit group 19  
     robotic control connection  
         control server cluster 15  
         redundant connections 15  
     SAN client 22  
     SAN connection  
         dynamic multi-pathing 15  
     SAN media server  
         application cluster 20  
     site  
         global cluster 22  
     storage device  
         redundant drives 16

**R**

replication considerations 57  
     clustering considerations 66  
     DNS considerations 58  
     master server considerations 57  
     networking considerations 58  
 restore data in cluster 74  
 restore full catalog 33  
 restore partial catalog 36  
 robotic control connection protection  
     control server cluster 15  
     redundant connections 15

**S**

SAN client protection 22  
 SAN connection protection  
     dynamic multi-pathing 15  
 SAN media server  
     non-failover cluster configuration 24  
 SAN media server protection  
     application cluster 20  
 site protection  
     global cluster 22  
 storage device protection  
     redundant drives 16  
 storage unit  
     configuring for robotic device 27

**U**

user backups 77  
 user-directed backups 73

**V**

virtual storage unit. *See* configuring with standalone  
     media server  
         upgrading media server 29