



Symantec

NetBackup Whitepaper

Simple site disaster recovery using "Recovery without import"

This document describes the process commonly known as "recovery without import", a technique which can be used to provide a simplified disaster recovery process for a NetBackup environment by making use of another production environment or a disaster recovery environment where the server configuration is significantly different to the production environment

If you have any feedback or questions about this document please email them to IMG-TPM-Requests@symantec.com stating the document title.

This document applies to NetBackup: 7.0, 7.1, 7.5 and all minor versions of these releases

This document is provided for informational purposes only. All warranties relating to the information in this document, either express or implied, are disclaimed to the maximum extent allowed by law. The information in this document is subject to change without notice. Copyright © 2012 Symantec Corporation. All rights reserved. Symantec, the Symantec logo and NetBackup are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners

Document Control

Revision History

Version	Date	Changes
1.0	10-Sep-2008	
2.0	30-Nov-2012	Updated for NetBackup 7.x

Related Documents

Version	Date	Title

Table of Contents

Introduction	1
Glossary	1
Process restrictions	1
Preparing a target domain for recovery without import	2
Using Partial Catalog Recovery	3
Changes to partial catalog recovery in NetBackup 7.5	3
Using Partial Catalog Replication	4
Changes to partial catalog replication in NetBackup 7.5	4
Operating recovery without import	4
Backing up in the DR domain	5
Recycling expired tapes	5

Introduction

This document describes the process, commonly known as “recovery without import”, which can be used to provide a simplified disaster recovery process for a NetBackup domain to another production domain or a disaster recovery domain where the server configuration is significantly different to the production environment.

The primary use of the recovery without import methodology is to allow backups to be recovered into another production NetBackup domain without the need to import the individual backup tapes. As such it is well suited to providing a disaster recovery capability where two production NetBackup domains exist.

The process works by simply presenting the image component of the source NetBackup catalog to the target Master Server. The tape records associated with the images must also exist in the target server’s EMM database but because they will not be used for backups they can simply be associated with a pool that is not used by backups in the target domain (NetBackup does not rely on media pool information when restoring backups).

Tapes are added into a private media pool so that they cannot be used or overwritten in the target domain. This allows backups written to tapes in the original domain to be restored in the new domain without needing to import the individual tapes.

This procedure assumes that the receiving site has a fully functioning NetBackup Master Server and there is sufficient bandwidth to copy the NetBackup catalog files from one site to another. If this procedure is being used as part of an overall disaster recovery plan, there are many other steps that are required and which must be addressed.

This process is intended for large scale recovery and has some inherent limitations in respect of media re-use. When attempting to recover individual backups in a different domain it is recommended that the import method is used rather than this approach.

Note: A significant potential for accidental data loss exists if the tape media are not physically write protected in the disaster recovery environment as they are not assigned in NetBackup and can therefore be overwritten if extreme care is not taken to protect against this.

Glossary

The following terms are used throughout this document:

Source domain – this is the NetBackup domain in which the backups were run and the backup image records were originally created. In a disaster recovery scenario this would be the production domain.

Target domain – this is the NetBackup domain to which the backup image records and the associated tapes are transferred prior to attempting recovery. In a disaster recovery scenario this would be the disaster recovery domain. Note that the disaster recovery domain may also be a production domain in its own right.

Private pool – a media pool in the target domain to which all media records from the source domain are added. This media pool must not be used by any backup policies in the target domain. As the records are created in this pool directly the media do not get returned to the scratch pool even though they are unassigned and cannot be overwritten as no policy in the target domain uses this pool.

Note: To ensure that tapes are not overwritten extreme care must be taken to ensure that the private pool **is not** used as a backup destination by any backup policy or Storage Lifecycle Policy. It is strongly recommended to write protect physical media wherever possible.

Process restrictions

The following restrictions apply to the recovery without import process:

1. The process can only be used for backups written to tape or BasicDisk storage.

2. All tapes belonging to the source and target domains must have unique labels and bar codes (i.e. no tape that exists in the source domain can exist in the target domain prior to the transition).
3. The media types and densities associated with the tapes must be the same in both domains (e.g. the process will not work if HCART3 is used for LTO3 in one domain and T10000 in the other domain).
4. The names of the backup clients in the source and target domains must be unique (i.e. no client that exists in the source domain can exist in the target domain prior to the transition).
5. Suitable tape libraries and drives exist at the target location to support restore from the tapes transferred from the source location.
6. Tapes from the source domain must be added directly to a pool in the target domain that is not used for any backups in that domain – there are no specific naming conventions for this pool but it is recommended that the user selects a name which makes the pool's function obvious.
7. Both Master Servers must be similar architecture (e.g. either UNIX/Linux or Windows).
8. Any “in progress” Storage Lifecycle Policy operations will not complete in the target domain.
9. Disk staging storage units will function as simple BasicDisk storage units in the target domain.
10. Tapes transferred from one domain to another are not automatically moved to the scratch pool when the backups on them expire and must be manually moved to the scratch pool.

Preparing a target domain for recovery without import

It is possible to prepare a target domain for recovery without import so that it can quickly be brought into operation when required. The following steps can be carried out at any time to ensure that recovery without import can be used:

1. Create a ‘private’ media pool in the target domain to house tapes sent from the source domain. This pool must not be used by any backup policies in the target domain.
Note: To ensure that tapes are not overwritten extreme care must be taken to ensure that the private pool **is not** used as a backup destination by any backup policy or Storage Lifecycle Policy. It is strongly recommended to write protect physical media wherever possible.
2. Add all the media from the source domain to this pool. This can be done using the CLI or GUI functions to simply add a range of tapes specifying the particular pool. Where the tapes to be sent from the source domain to the target domain form a specific set or range then only these tapes need to be added to the target domain.
3. Create alternate read host entries in the target domain for each of the Media Servers in the source domain (FORCE_RESTORE_MEDIA_SERVER entries in bp.conf on UNIX and Linux and alternate restore server entries in the Master Server host properties ‘server’ tab in Windows).
4. If desired, set up the replication of the image database from the source domain to the target domain and create the necessary soft links or ALTPATH entries on the target Master Server to allow the image records to be accessed when the replicated storage is presented to it.
5. Set up a process for transferring tapes between the source and target domains on a daily basis (e.g. implement the NetBackup Vault Option). If replication is not being used this should include a backup of the image database which should be clearly identified for tracking purposes.

Where two production domains are operating as targets for each other these steps should be carried out in both domains.

Once implemented, this configuration can be left in place until it is required for disaster recovery purposes place.

Using Partial Catalog Recovery

Partial catalog recovery allows the target domain to be populated with image data from the source domain by restoring it from the source domain's catalog backup. Because only the image information is recovered from the source domain the device and server configuration of the target domain is unaffected.

Recover the catalog backup taking the option to recover only the catalog image and configurations files (the P option when running `bprecover -wizard`).

Where two domains are to be used to provide recovery positions for each other it is a simple matter of adding all the tapes from one domain to a single pool (perhaps using the name of the Master Server) on the other domain and configuring alternate restore server settings for each Media Server in the alternate domain.

If BasicDisk storage is used in the source domain this storage may be replicated to the target domain and mounted against an alternate Media Server, however the same mount point must be used on the alternate Media Server to allow restore operations.

Note: BasicDisk storage which functions as staging storage in the source domain will not function as staging storage in the target domain and images that have not been duplicated to tape at in the source domain must be manually duplicated to tape in the target domain (note that the images will not be removed from disk even when the images expire).

If Enterprise Disk storage is used in the source domain this storage may also be replicated “out of band” to the target domain. Replicated Enterprise disk storage must be synchronized in the target domain using the `nbcatsync` command once partial catalog recovery is complete. Refer to [TECH66060](#) for more details of this process.

Note: Any “in progress” storage lifecycle policy operations from the source domain will not complete in the target domain because information about active operations is stored in the relational database and is not available in the target domain.

Changes to partial catalog recovery in NetBackup 7.5

One of the significant changes made in NetBackup 7.5 was to move the image metadata (the header files from the image database) into the relational database. This change was made both to improve performance and scalability and to allow for the introduction of new fields within the headers to support new features planned for future releases.

This change has made it necessary to change the way in which partial catalog recovery works and exactly what gets recovered. This, in turn, may mean that existing disaster recovery procedures based on partial catalog recovery will need to be modified after upgrading to NetBackup 7.5.

Prior to NetBackup 7.5 the partial catalog recovery option simply restores the complete flat file component of the catalog along with the license database and, in the case of UNIX and Linux environments, the `bp.conf` file. The relational database components are not restored.

Starting with NetBackup 7.5 the default behavior of the partial catalog recovery option does two things:

1. Restores the image database component of the flat file database (`/usr/opensv/netbackup/db/images` or the Windows equivalent) to its original location.
2. Restores the relational database to the designated relational database staging area.

In order to complete the partial catalog recovery the image metadata must be exported from the copy of the relational database in the staging area and imported into the live relational database. This is achieved by running the following commands:

1. `cat_export -all -staging -source_master <source master server>`
2. `cat_import -all -replace_destination`

Using Partial Catalog Replication

Partial catalog replication avoids the need to perform a catalog recovery operation in the target domain and thus saves time during the recovery process. It is, however, more difficult to configure and operate than the partial catalog recovery method described in the previous section. Replicate the image component of the source catalog to the target Master Server and present the replicated data to the target Master Server. In this case the presentation will vary depending on the platform used.

- a. For UNIX and Linux servers the replicated volume is mounted alongside the catalog volume and soft links are created at the client level on the catalog volume
- b. For Windows the replicated volume is mounted alongside the volume containing the catalog and ALTPATH entries are created at the client level on the catalog volume. Note that in order to replicate only the image component on Windows the source catalog must also make use of ALTPATH settings to map the images to a separate volume so that they can be replicated.

Note that where replication is used to transfer image records sufficient bandwidth must exist between the sites to allow the replicated data to be updated without impacting running backups on the source site.

As in the case of partial catalog recovery, disk storage devices can be replicated to the target domain. However a catalog backup must exist in the source domain and the image record must be replicated to the target domain to enable the `nbcatsync` command to work.

Changes to partial catalog replication in NetBackup 7.5

The changes to the NetBackup catalog structure in NetBackup 7.5 have a significant impact on how partial catalog replication is implemented. In order to ensure that the catalog metadata is available in the target domain the relational database components of the source domain must be regularly backed up to the replicated storage. The steps required to implement and operate partial catalog replication with NetBackup 7.5 and above are discussed in more details in [TECH200290](#)

Operating recovery without import

If the steps outlined in the previous section are followed then operating recovery without import simply involves the following steps:

1. Present the source domain image information to the target domain by either mounting the replicated storage or restoring the backup of the image database to the target Master Server.
2. If the tapes sent to the target domain are duplicate (copy 2) tapes then one of the following approaches must be used to ensure that they are selected by restore jobs:
 - a. Create the file `ALT_RESTORE_COPY_NUMBER` in the NetBackup root directory (`/usr/opensv/netbackup` or `<install path>\netbackup`) containing the copy number to be used for restores – this value is then applied to all restores for all clients until the file is removed.
 - b. Add the qualifier `-copy 2` to the CLI restore command `bprestore`
 - c. Promote the duplicate tapes to be primary copy tapes by running the following command for each client you need to restore:

```
bpchangeprimary -copy 2 -cl <client name>
```

Note that this must be done after the latest image information has been presented to the target Master Server.

3. Start restoring the backups.

Backing up in the DR domain

Once the backups have been restored at the DR site and production has resumed it will probably be necessary to start running backups at that site. Backup policies may be re-activated but care must be taken to ensure that the policies do not use the private pool and that a suitable stock of tapes exists in the global scratch pool at the DR site.

Recycling expired tapes

As previously mentioned tapes which have been moved to a private pool under this method are not assigned within NetBackup and thus do not expire and get re-cycled to the scratch pool when all the backups on them have expired. To ensure that sufficient tape stocks are maintained at the DR site a periodic check of the databases should be carried out and tapes that have no backups on them should be moved to the scratch pool.

The simplest way of checking this is to create two lists by running the commands `bpimagelist -hoursago 100000 -media -l` and `vmquery -pn <private pool name> -b` and 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>`.

About Symantec:

Symantec is a global leader in providing storage, security and systems management solutions to help consumers and organizations secure and manage their information-driven world.

Our software and services protect against more risks at more points, more completely and efficiently, enabling confidence wherever information is used or stored.

For specific country offices and contact numbers, please visit our Web site: **www.symantec.com**

Symantec Corporation
World Headquarters
350 Ellis Street
Mountain View, CA 94043 USA
+1 (650) 527 8000
+1 (800) 721 3934

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.