NetBackup™ in Highly Available Environments Administrator's Guide

Windows, UNIX, and Linux

Release 11.2

VERITAS
now part of Cohesity

NetBackup™ in Highly Available Environments Administrator's Guide

Last updated: 2026-05-28

Legal Notice

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to Cohesity Support, to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the Cohesity Support Portal to create a new case.
- Click the (**?**) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the Cohesity Knowledge Base.
- Log in to the Cohesity Support Portal to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

   Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

Chapter 1

About in this guide

This chapter includes the following topics:

What's in this guide

The *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 primary servers. In addition, the guide details operating practices when catalogs are replicated between clustered or non-clustered NetBackup primary 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 primary server clustering, refer *NetBackup Clustered Primary Server Administrator's Guide*.

https://support.cohesity.com/s/article/article-100040135

See “Documents related to NetBackup in highly available environments” on page 7.

Documents related to NetBackup in highly available environments

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

- Refer *NetBackup Clustered Primary Server Administrator's Guide* for information about clustering NetBackup.
  https://support.cohesity.com/s/article/article-100040135
- Refer *NetBackup Installation Guide* for information about installing NetBackup.
  https://support.cohesity.com/s/article/article-100040135
- Refer *NetBackup Administrator's Guide, Volume I and Volume II*, for general information about NetBackup.
  https://support.cohesity.com/s/article/article-100040135

Chapter 2

NetBackup protection against single points of failure

This chapter includes the following topics:

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 "Network link failures" on page 10.
Storage device connections	See "Storage device connection failures" on page 10.
Storage devices	See "Storage device failure" on page 11.
Media availability	See "Media availability failures" on page 11.

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

Point of failure	Protection method
Primary server	See "Primary server failures" on page 12.
Media server	See "Media server failures" on page 13.
LAN client	See "LAN client failures" on page 17.
SAN client	See "SAN client failures" on page 17.

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

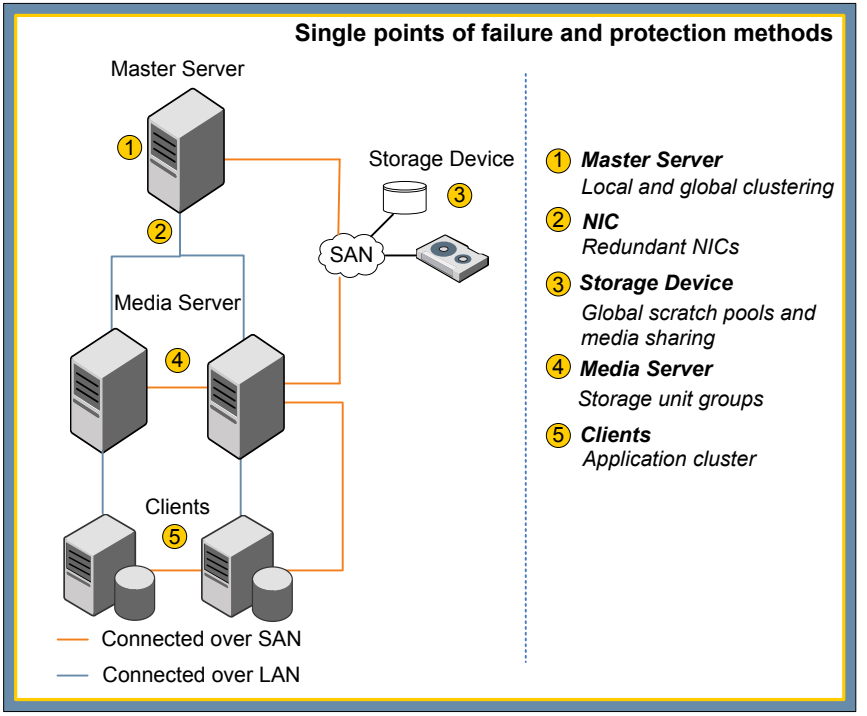

Network link failures

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.

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 "SAN connection failures" on page 10.

See "Robotic control connection failures" on page 10.

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.

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 primary 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 primary 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.

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.

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	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. 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. 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.
Media sharing	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. To enable media sharing, set the **Volume Pool** properties to use the **Maximum number of partially full media** 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.

Primary server failures

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

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

https://support.cohesity.com/s/article/article-100040135

The primary servers that are running in virtual machines can be protected using the Hypervisor's high availability tools. For details refer to http://www.veritas.com/docs/000006177.

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 "Dedicated media server failures" on page 13.
Non-dedicated media servers	Run other applications also that require backing up. Also back up data from other systems. See "Non-dedicated media servers failures" on page 14.
SAN media servers	Run other applications also that require backing up. Do not back up data from other systems. See "SAN media server failures" on page 15.

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.

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.

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.

Restoring tape backups using an alternative media server

When you restore 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 main 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 (Windows)

1. Sign in to the primary server.
2. On the left, click **Hosts > Host properties**.
3. Select the primary server and click **Connect**.
4. Select the primary server and click **Edit primary server**.
5. Click **Restore failover**.
6. Click **Add** to add a media server.

To configure the failover restore media servers (UNIX)

1. Sign in to the primary server.
2. In the `bp.conf` file, create an entry `FAILOVER_RESTORE_MEDIA_SERVER`.

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.

If you encounter this problem, configure the force restore media server setting, with one of the following methods.

Configure the force restore media server setting

To configure the failover restore media server in the `bp.conf` file (UNIX)

1. Sign in to the primary server.
2. (UNIX) In the `bp.conf` file, create an entry `FAILOVER_RESTORE_MEDIA_SERVER`.

To add a media host override (Windows)

1. Sign in to the primary server.
2. On the left, click **Hosts > Host properties**.
3. Select the primary server and click **Connect**.
4. Select the primary server and click **Edit primary server**.
5. Click **General server**.
6. Locate **Media host override**. Then click **Add** to add a media 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`

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.

Create the touch file as follows:

- (Linux) On the primary server, create the file in `/usr/openv/netbackup/db/config`
- (Windows) On the primary server, create the file in `<install path>\veritas\netbackup\db\config`.
  `USE_BACKUP_MEDIA_SERVER_FOR_RESTORE` is a global setting and always forces the restore to the server that did the backup.

Use the `bprestore -disk_media_server` command

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.

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.

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.

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 primary 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 alternate 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 *NetBackup Administrator's Guide, Volume I*.

https://support.cohesity.com/s/article/article-100040135

Protecting the 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 18. discusses the various methods for protecting NetBackup catalogs.

Table 2-5 NetBackup catalog protection in highly available environments

Protection Method	Description
Catalog backups	The catalog backup protects the NetBackup catalog on the primary 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. Catalog backups written to tape use media from the Catalog Backup volume pool only. For more information, refer to the *NetBackup Administrator's Guide, Volume I*.

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

Protection Method	Description
Catalog replication	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. Replicating the catalog to a standby primary server at the disaster recovery or alternate 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. **Note:** Replication does not protect against catalog corruption or accidentally deleting or expiring images. You must make regular scheduled catalog backups. See "About NetBackup catalog replication" on page 32. See "About catalog recovery" on page 21.

Chapter 3

About site disaster recovery with catalog backup and recovery

This chapter includes the following topics:

Disaster recovery packages

For increased security, a disaster recovery package is created during each catalog backup. The disaster recovery package file has `.drpkg` extension.

The disaster recovery (DR) package stores the identity of the primary server host. NetBackup requires this package to get the identity of the primary server back after a disaster. Once you have recovered the host identity, you can perform the catalog recovery.

The disaster recovery package contains the following information:

- NetBackup CA-signed certificates and private keys of the primary server certificate and the NetBackup certificate authority (CA) certificate
- Information about the hosts in the domain
- Security settings
- External CA-signed certificates
  External CA-signed certificates from Windows certificate store, if applicable

- NetBackup configuration options that are specific to external CA-signed certificates
- Key management service (KMS) configuration

Note: By default, the KMS configuration is not backed up during catalog backup. Set the KMS_CONFIG_IN_CATALOG_BKUP configuration option to 1 to include the KMS configuration as part of the disaster recovery package during catalog backup.

Note: You must set a passphrase for the disaster recovery package for the catalog backups to be successful.

About catalog recovery

A major problem that 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 primary server. NetBackup catalog backups are primarily intended for recovering from catalog storage or primary server failure rather than site loss. The default scenario is that NetBackup restores the complete catalog including the NetBackup database. Primary servers use the catalog information to direct backups and restores, query the media servers, and establish the status of the backup devices. In a DR environment which does not contain these media servers, the performance of the primary server can be affected. Also, the ability to carry out restore operations can be affected, as polling operations fail to connect and time out.

Note: In a cluster setup, if you use external CA-signed certificates for host communication, ensure that the virtual name and the cluster nodes have the same certificate authority (CA) usage. For example, if a node uses only external CA-signed certificates, ensure that the virtual name also uses external CA-signed certificates. If there is a mismatch in the CA usage of the virtual name and cluster nodes, catalog backup and catalog recovery may fail.

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 22.

- In the partial catalog recovery the NetBackup database is not restored.
  See "About partial catalog recovery" on page 26.

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 39.
- See "Performing full catalog restore" on page 23.
- See "Performing a partial catalog restore" on page 27.

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 database components, which include the storage unit definitions, media assignment, and history.
- It retains the tape information from the main site including the media pool and other assignment information.
- It restores the NBDB, NBAZDB, and BMR (if configured) 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:

- Catalog recovery does not recover host certificates. To recover the NetBackup primary server identity or host certificates and other information, the disaster recovery package must be recovered.
  See "Disaster recovery packages" on page 20.
- When you recover the 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 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 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 primary 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 NetBackup 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. Run the `nbgetconfig` command and save the output. This output can be used after the catalog recovery to recover the host-specific information that is overwritten during the catalog recovery.

   For example:

   ```
   ./nbgetconfig > sample.txt
   ```

2. Run the `bprecover` command to recover the entire catalog.

   Note: The DR primary server must have the same name and topology as the production primary server. If the production primary server is a cluster then the DR primary 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. After you run the `bprecover` command, set a passphrase for the disaster recovery package so that subsequent catalog backups are successful.

   See “Disaster recovery packages” on page 20.

4 During catalog recovery, security certificates for cluster nodes are not recovered. Only the virtual name certificate is recovered.

If NetBackup certificates are used for host communication	For successful host communication, you must deploy NetBackup certificates (host name-based and host ID-based certificates) on all cluster nodes after a disaster. For more details, refer to the *Generating a certificate on a clustered primary server after disaster recovery installation* chapter from the NetBackup Security and Encryption Guide.
If external certificates are used for host communication	For successful host communication, you must configure all cluster nodes to use external certificates after a disaster. For more details, refer to the NetBackup Security and Encryption Guide.

5 Clear allowed list cache and restart the service on all hosts in the domain.

6 Deactivate all the backup policies to prevent backups from starting automatically. Use one of the following methods:

- The NetBackup web UI.
- Run the `bpplinfo <policy> -modify -inactive` CLI.

7 Shut down NetBackup.

8 Recover the host settings that you backed up in step 1. Run the following command.

```
./nbsetconfig sample.txt
```

9 Start the NetBackup Scale-Out Relational Database Manager, NetBackup PBX, and EMM services on the new primary server.

- On Linux primary servers, run the following commands:
  - `/usr/openv/netbackup/bin/nbdbms_start_stop start`
  - `start /opt/VRTSpbx/bin/pbx_exchange`
  - `/usr/openv/netbackup/bin/nbemm`
- On Windows primary servers, start the following Windows services:
  - NetBackup Scale-Out Relational Database Manager
  - Cohesity 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 Scale-Out Relational Database Manager, see the NetBackup Troubleshooting Guide.

10 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 primary_server
```

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

```
nbemmcmd -deletealldevices -allrecords
```

12 This step is required if you have NAT clients in your environment.

If the NetBackup Messaging Broker (or `nbmqbroker`) service is configured, after the catalog restore you need to enable the cluster to monitor the service using the `configureMQ -enableCluster` command.

For more information on the command, refer to the NetBackup Command Reference Guide.

13 Restart NetBackup.

14 Create the new tape drive and library configuration.

15 If the barcode masking rules were used at step 9, ensure that the same rules are set here. If necessary, add them.

16 Verify that all the recovery media are set to non-robotic.

17
- 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.

18 Once all the recovery media are set to non-robotic, in the **Inventory all the tape libraries** 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.

If you have configured an external CA-signed certificate for the NetBackup web server, you must run the `configureWebServerCerts` command on the active node. This action ensure that the external certificate is used after the failover.

For more information on the command, see the NetBackup Commands Reference Guide.

For all cluster nodes, do the following:

- Define the external certificate configuration options (`ECA_CERT_PATH`, `ECA_CRL_PATH`, and so on) in the configuration file on the node.
- Run the `nbcertcmd -enrollCertificate` on the node.
  For more details, refer to the *NetBackup Security and Encryption Guide*.

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. If the catalog recovery is not performed immediately after the DR package recovery, to make the DR environment consistent, do the following:
2. Modify the backup policies, including the catalog backup policy, to use the storage units available at the DR site and enable them.
3. Delete the backup policies that are no longer required.
4. Delete the storage units that are associated with the media servers and are not part of the DR environment.
5. Modify any Storage Lifecycle Policies that use storage units that you have deleted.

About partial catalog recovery

Auto Image Replication (A.I.R.) is recommended to duplicate backups across multi-domain configurations. Partial catalog recovery may be used for this purpose in situations where A.I.R. is not an option.

Partial catalog recovery recovers only the flat file components and not the database. This way the details of the existing infrastructure (servers, devices etc.) at the DR site are 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 servers 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 primary server topology at the DR site does not need to reflect the topology at the production site. Thus, a catalog backup from a clustered primary server can be restored to a standalone primary 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 main 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 a 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. Run the `nbgetconfig` command and save the output. This output can be used after the catalog recovery to recover the host-specific information that is overwritten during the catalog recovery.

   For example:

   ```
   ./nbgetconfig > sample.txt
   ```

2. Recover only the NetBackup catalog image and configuration files.
   - Select the **Partial catalog recovery** option when prompted.
   - Or run the `bprecover -wizard` command. When you are prompted `"Do you want to recover the entire NetBackup catalog? (Y/N)"`, enter **N**.

Note: The DR primary server must have the same name as the production primary 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 If you want to recover the image header information without recovering the entire NetBackup database, perform the following steps:

- Step a - Back up the target database. Run the following command.

  `nbdb_backup -online` *`directory`*

  Make sure that you do not specify the staging folder as the output directory. (The staging folder contains the schema data and configuration data for the NetBackup database from the catalog backup. Image `.f` and configuration files are recovered to their final destinations.)

- Step b - Recover the NetBackup database from the staging directory.

  ```
  nbdb_restore -recover -staging
  ```

- Step c - Export the image header data that you want to import from the backup.

  For example, the following command exports export all image header data. The data is exported to the `netbackup/db.export` directory.

  ```
  cat_export -all
  ```

- Step d- Recover the NetBackup database with the following command.

  `nbdb_restore -recover` *`directory`*

  Make sure that you specify the same directory as in step a.

- Step e- Run the `cat_import` command to import the image header data that you extracted in step c.

  ```
  cat_import -all -replace_destination -delete_source
  ```

  The command does the following:

  - Imports all of the image header data in the `netbackup/db.export` directory.
  - Replaces any image header data that was exported that already exists in the target database.
  - Removes the image header data that resides in the `netbackup/db.export` directory.

- Step f- If you recovered the catalog from a disk device, you may have to fix the disk media ID references. Run the following command:

  ```
  nbcatsync -sync_dr_file DR file path -dryrun
  ```

  Replace *DR file path* with the path to the catalog DR file.

- Step g - If the result of the dry run is satisfactory, run the following command:

  ```
  nbcatsync -sync_dr_file DR file path
  ```

4. Deactivate all the backup policies to prevent backups from starting automatically, in one of the following ways.
   - In the NetBackup web UI.
   - Run the `bpplinfo <policy> -modify -inactive` CLI.
5. Shut down NetBackup.
6. Recover the host settings that you backed up in step 1. Run the following command.

   ```
   ./nbsetconfig sample.txt
   ```

7. Start NetBackup.
8. 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 primary 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 the 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.

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 primary server platforms. You can use it with all AdvancedDisk types supported by NetBackup.

To recover a disk in a cross-domain replication DR environment

1. Sign in to the DR domain's primary server.
2. 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. Run the following command:

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

3. Perform a partial catalog recovery from the replicated catalog backup. Run the following command:

   ```
   bprecover -wizard
   ```

   When you are prompted `"Do you want to recover the entire NetBackup catalog? (Y/N)"`, enter **N**.

4. Run the `cat_export -all` to export the metadata from the replicated database backup.
5. Run the command `cat_import -all` to import the exported metadata into the active database.
6. 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. Run the following command:

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

Chapter 4

About site loss protection with auto image and catalog replication

This chapter includes the following topics:

About Auto Image Replication (AIR)

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. Cohesity 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 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.

Dynamic multi-streaming support of catalog image for AIR DR scenarios does not ensure image grouping completeness on the target domain. This is because, the bulk/batch operations should ensure consistency of catalog images.

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 35.
- Multi-site cross domain replication
  See "About multi-site cross domain replication" on page 38.

About conditions for support of replicated NetBackup catalogs

A NetBackup environment that is 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 main primary 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 main NetBackup server can get corrupted due to replication to the alternate 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. Cohesity 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 primary server must reside on the same virtual server that is controlled as a single entity.

- The primary and the alternate primary servers must be of similar type, specification, operating system, and use the same virtual host name.
- The alternate primary server must not have any other NetBackup function, neither in the same domain as the main primary server, nor in another domain. For example, you cannot use the alternate primary server as a media server if it is not used as a primary server. You also cannot use it as a primary server for another NetBackup domain. Catalogs are replicated but cannot be merged.
- Configure both the clustered and the non-clustered environments to use a virtual host name and IP address for the NetBackup primary server that is separate from the physical host names and IP addresses of the servers. Separate virtual host name and IP address let you control the active primary server node through DNS routing. It also prevents the primary and the alternate primary 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 host name must be specified during installation.
- Ensure that the main primary server and the alternate primary 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 primary 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 main and the alternate 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 primary servers the cluster common volume is replicated.

  For non-clustered primary servers, for details of the paths that must be linked to a volume set for replication,
- Ensure that the virtual name or DNS alias does not resolve to both the main and the alternate hosts at the same time.
- Catalog replication does not remove the requirement for catalog backup. Regularly back up the NetBackup catalog from the main primary server to protect against accidental image expiration or other inconsistencies that are introduced in the catalog on the main site and replicated to the alternate 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 an alternate primary server lets you restore data during a short-term outage of the main primary 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 main 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 alternate site. This usage is not a requirement for support.
- Both the catalog and the backup images must be accessible at the alternate 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 alternate 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.
- Microsoft Distributed File System Replication (DFSR) technology is not supported as it does not guarantee write-ordered consistency of the files being replicated. For more information, see https://support.cohesity.com/s/article/article-100043283

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 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 primary 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 alternate primary server.

In the multi-site single domain model, NetBackup catalogs are replicated between the sites. In the event of a problem at the main site, the primary server is failed over to a standby node on the alternate 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 primary server clustering and storage replication. This combination allows the primary server to be relocated easily and quickly to the alternate 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 36.
- Multi-site single domain with optimized duplication
  See “About multi-site single domain with optimized duplication” on page 37.

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

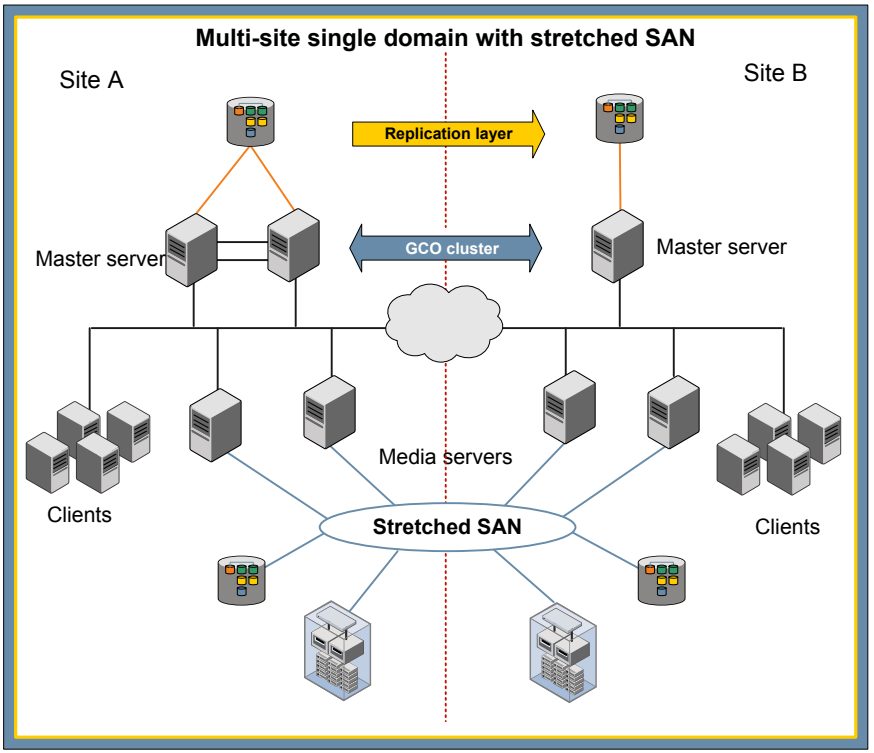

About multi-site single domain with optimized duplication

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

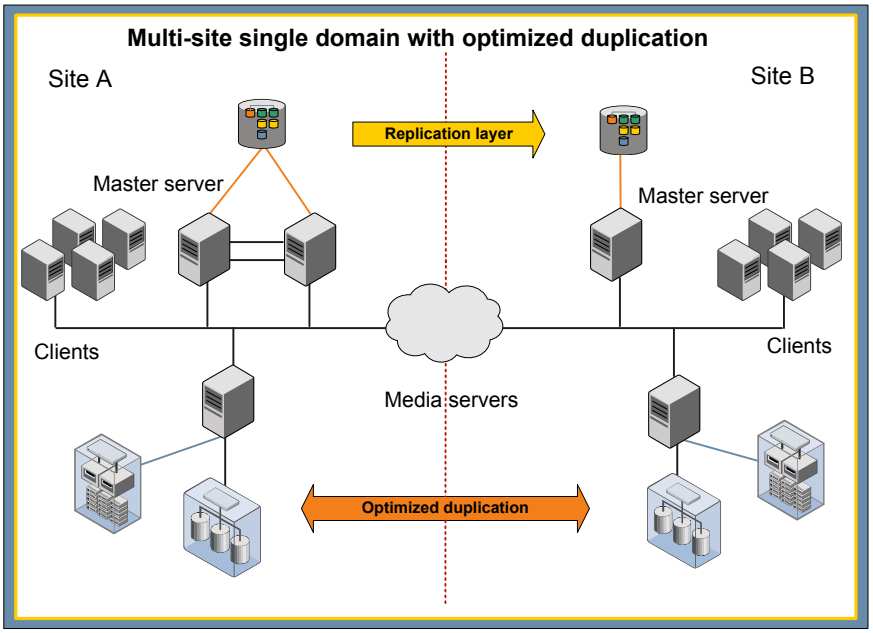

About multi-site cross domain replication

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 primary 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 alternate primary server in the DR domain, ensure that the primary server, media servers, network connections, and NetBackup software are functional.

Cohesity 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. Refer to the following steps when you prepare a DR plan:

To plan a cross domain replication disaster recovery domain

1. Install the same NetBackup version on the primary server, media servers, and clients in the DR domain that is used at the production domain.

   Note: 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 primary server and the media servers in the DR domain.

   If the full catalog replication method is used and the primary server at the production domain is clustered, a clustered primary 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 primary server in the DR domain is not required.

   ---

2. Test the network connectivity and authentication between the clients and servers using test backup policies. Disable the policies after testing.
3. 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.
4. Set the FAILOVER_RESTORE_MEDIA_SERVER 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.

5. If the partial replication method is used, create a non-scratch Media Pool which is not used by any backup policy. Configure barcode rules to ensure that the backup tapes are automatically added to that pool.
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.
7. Ensure the following:
   - If the original backup tapes are used for DR purposes, they must be loaded in the tape libraries in the DR domain.
   - 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.

Note: Cohesity 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.

About full catalog replication

In full catalog replication, all parts of the catalog are replicated to the alternate primary 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 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 primary 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 NetBackup database.

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

To recover the catalog with full catalog replication

1. Ensure that replication between the main and the alternate sites is stopped.

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

2. Mount the replicated volume to the appropriate mount point on the alternate primary server.

3. Start the NetBackup Scale-Out Relational Database Manager, NetBackup PBX, and EMM services on the new primary server.

   - On Linux primary servers run the following commands:
     - `/usr/openv/netbackup/bin/nbdbms_start_stop start`
     - `/opt/VRTSpbx/bin/pbx_exchange`
     - `"/usr/openv/netbackup/bin/nbemm -maintenance`
   - On Windows primary servers start the following Windows services:
     - NetBackup Scale-Out Relational Database Manager
     - Cohesity 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 primary_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, deactivate all the backup policies to prevent backups from starting automatically. Use one of the following methods:

- The NetBackup web UI.
- Run the `bpplinfo <policy> -modify -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 Verify that 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 tape libraries** 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 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 alternate primary server.

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

Preparing an environment for partial catalog replication

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

To prepare an environment for partial catalog replication

1. Change the configuration on the source (production) primary server to ensure that the staging area for the 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. Back up the database several times per day (ideally each hour) by running the following command in a scheduled script.

   ```
   nbdb_backup -online directory
   ```

Recovering the environment with partial catalog replication

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

To recover the environment with partial catalog replication

1. Run the `nbgetconfig` command and save the output. This output can be used after the catalog recovery to recover the host-specific information that is overwritten during the catalog recovery.

   For example:

   ```
   ./nbgetconfig > sample.txt
   ```

2. Ensure that replication is stopped between the main and the alternate sites.

   Replication stops if the main primary server is unavailable or if the replication link is disabled.

3. Mount the replicated volume to the appropriate mount point on the alternate primary server.

4. Use the command `nbdb_admin –vxdbms_nb_staging <directory>` on the target (disaster recovery) primary server to point the staging area for the database to the location on the replicated storage.

5. Recover the image header information without recovering the entire NetBackup database, perform the following steps:

   - Step a - Back up the target database. Run the following command.

     `nbdb_backup -online` *`directory`*

     Make sure that you do not specify the staging folder as the output directory. (The staging folder contains the schema data and configuration data for the NetBackup database from the catalog backup. Image `.f` and configuration files are recovered to their final destinations.)

   - Step b - Recover the NetBackup database from the staging directory.

     ```
     nbdb_restore -recover -staging
     ```

   - Step c - Export the image header data that you want to import from the backup.

     For example, the following command exports export all image header data. The data is exported to the `netbackup/db.export` directory.

     ```
     cat_export -all
     ```

   - Step d- Recover the NetBackup database with the following command.

     `nbdb_restore -recover` *`directory`*

     Make sure that you specify the same directory as in step a.

- Step e- Run the `cat_import` command to import the image header data that you extracted in step c.

  ```
  cat_import -all -replace_destination -delete_source
  ```

  The command does the following:

  - Imports all of the image header data in the `netbackup/db.export` directory.
  - Replaces any image header data that was exported that already exists in the target database.
  - Removes the image header data that resides in the `netbackup/db.export` directory.

- Step f- If you recovered the catalog from a disk device, you may have to fix the disk media ID references. Run the following command:

  ```
  nbcatsync -sync_dr_file DR file path -dryrun
  ```

  Replace *`DR file path`* with the path to the catalog DR file.

- Step g - If the result of the dry run is satisfactory, run the following command:

  ```
  nbcatsync -sync_dr_file DR file path
  ```

6. Run the command `cat_import -all` to import the exported metadata into the active database.
7. Recover the host settings that you backed up in step 1. Run the following command.

   ```
   ./nbsetconfig sample.txt
   ```

8. Start NetBackup on the secondary primary server.
9. If the backup policies are replicated, deactivate all backup policies to prevent backups from starting automatically. Use one of the following methods:
   - The NetBackup web UI.
   - Run the command `bppllist <policy> -set -inactive`.
10. Ensure that the appropriate FAILOVER_RESTORE_MEDIA_SERVER settings are defined to direct restore operations through the media servers at the alternate site.

11 To restore the backups from tapes, you must add the tapes to the disaster recovery primary server's catalog. Place the tapes in a tape library and run an inventory of the library. To prevent the tapes from being accidentally overwritten the disaster recovery primary server should have a barcode rule that adds the tapes to a volume pool. The volume pool must not use the global scratch pool and no backup policies should use this pool. Ideally the tapes should also be physically write locked.

12 For disk-based backups, the storage servers and disk pools must be added to the disaster recovery primary 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. After 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 primary 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>`.

Chapter 5

Deploying NetBackup primary servers with full catalog replication

This chapter includes the following topics:

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
Primary server considerations	Cohesity does not recommend operating a primary server as a combined primary 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. Catalog replication is not a substitute for catalog backup and the catalog must be backed up on a regular basis.

Table 5-1 Replication considerations *(continued)*

Considerations	Description
Networking considerations	In a multi-site single domain configuration, the primary 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. It is recommended that a link of at least 10 Mb/sec and ideally 100 Mb/sec must be provided between the sites to handle this traffic. If host-based replication is used, additional I/P bandwidth is required for the replication layer. The additional bandwidth must also be factored in.
DNS considerations	If the primary server nodes at the alternate site are on a different subnet from the primary server nodes at the main 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. **Note:** To propagate the DNS change automatically by the cluster service group, the DNS resource must come online after starting NetBackup.
Main and alternate primary server considerations	In order to perform a failover when using catalog replication the main and alternate primary servers must use the same topology. The main and alternate site primary server nodes must both be either clustered or non-clustered. **Note:** The clustered primary servers do not require the same number of nodes at each site. For additional details refer to the following article: https://www.cohesity.com/docs/000090837

About non-clustered NetBackup primary server with catalog replication

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

Installing and configuring non-clustered NetBackup primary server with catalog replication

The installation and configuration of non-clustered NetBackup primary servers with catalog replication progresses through multiple stages which are described inTable 5-2

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

Table 5-2 Installing and configuring non-clustered NetBackup primary server with catalog replication

Step	Description
Stage 1	Make sure that you meet all the conditions of support before proceeding with the actual installation. See "About conditions for support of replicated NetBackup catalogs" on page 33.
Stage 2	Install and configure the non-clustered NetBackup primary server on the main site. See "Installing and configuring the main NetBackup primary server" on page 50.
Stage 3	Install and configure the non-clustered NetBackup primary server on the alternate site. See "Installing and configuring an alternate NetBackup primary server" on page 52.

Installing and configuring the main NetBackup primary server

The main primary server is the server that normally functions as the primary server. It must be installed first.

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

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

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

   Verify that the string is `NB_<alias name>` and not `NB_<hostname>` and modify as necessary.

   This file is located in the following directory:

   `<install path>\VERITAS\netbackupdb\data\vxdmbs.conf`

   `/usr/openv/db/data/vxdbms.conf`
5. Move the catalog components to the volume that is replicated to the alternate primary server.

   For Windows installations, map the following paths to a common volume. Use symbolic links.

   - `<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 Linux installations, soft link the following paths to locations on a common volume:

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

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

   To prevent this automatic start, make the changes in step 7 and step 8 after the initial installation and after applying patches or upgrades.

7 Make the following changes on the NetBackup primary server.

   - On a Linux primary server, remove the links to `/etc/init.d/netbackup` created during the installation to enable automatic startup.
     Refer to the NetBackup Installation Guide for details of links for each operating system.
   - On a Windows primary server, go to the **Services Manager** and set the **Startup type** for all the NetBackup services to **Manual**.

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

   At this stage, you can configure media servers and storage devices.

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

   On a Linux primary 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 a Windows primary 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 an alternate NetBackup primary server

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

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

1. Stop NetBackup on the main primary server.
2. Map the DNS alias name to the alternate primary server.
3. Install the NetBackup primary server on the alternate primary server node, specifying the alias name for the primary server. During installation, apply the same list of servers on the alternate primary server.
4. After the installation is complete, shut down NetBackup.
5. To ensure that NetBackup starts correctly when switching to the alternate primary server, verify the `NB_<alias name>` string in the `vxdbms.conf` file.

   Verify that the string is `NB_<alias name>` and not `NB_<hostname>` and modify as necessary.

   This file is located in the following directory:

   ```
   <install path>\VERITAS\netbackupdb\data\vxdmbs.conf
   ```

   ```
   /usr/openv/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.

   Note: During a failover operation the replicated volume is mounted on the alternate primary server and not the disk volume.

7. Move the catalog components to this disk volume.

   For Windows installations, map the following paths to a common volume. Use symbolic links.

   - `<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/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 alternate primary server. By default, NetBackup is started automatically when the primary server is started.

To prevent this automatic start, make the changes as per Steps 9 and 10 after the initial installation and after applying patches or upgrades.

9 Make the following changes on the NetBackup primary server.

- On a Linux primary server, remove the links to `/etc/init.d/netbackup` created during the installation to enable automatic startup.
  Refer to the NetBackup Installation Guide for details of links for each operating system.
- On a Windows primary server, go to the **Services Manager** and set the **Startup type** for all the NetBackup services to **Manual**.

10 Start NetBackup on the alternate primary server. Confirm that NetBackup comes up and then shut it down again. At this stage, you can configure media servers and storage devices.

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

On a Linux primary 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 primary server:

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

11 After NetBackup shuts down, dismount the disk volume that is mounted on the primary server (refer step 6). Then reset the DNS alias name to the main primary server. Then restart NetBackup on the main primary server.

Upgrading NetBackup primary server in a non-clustered replicated configuration

In order for global failover to work correctly, both the main and the alternate site clusters must run the same version of NetBackup. This means that both clusters must be upgraded at the same time. The upgrade process requires the replication link to be disabled and each cluster to be upgraded independently.

To upgrade the NetBackup primary server in a non-clustered replicated configuration

1. Suspend replication between the main and the alternate sites.
2. Upgrade NetBackup on the main primary server by following the standard upgrade procedures. (Refer to the NetBackup Upgrade Guide.)
3. Run backup and restore tests to confirm that the upgrade was successful.
4. Ensure that the alternate primary server is isolated from the wider network so that it cannot contact the media servers and clients when it comes online.
5. Bring the alternate primary server online and mount the replicated catalog volume to it.
6. Upgrade NetBackup on the alternate site primary server following the standard upgrade procedures. (Refer to the NetBackup Upgrade Guide.)
7. After the upgrade is complete take the alternate site primary server offline. Doing this action avoids unnecessary and potentially time consuming post processing operations on the alternate site catalog volume.
8. Re-connect the alternate site primary server to the wider network.
9. Restart the replication process and allow the replicated volume to fully synchronize.

About globally clustered NetBackup primary servers with catalog replication

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

Installing and configuring a globally clustered NetBackup primary server with catalog replication

The installation and configuration of clustered NetBackup primary servers with catalog replication progresses through multiple stages which are described in Table 5-3.

Table 5-3 Installing and configuring clustered NetBackup primary server cluster with catalog replication

Stage	Description	Action
Stage 1	Installation prerequisites	Make sure that you meet all the conditions of support before proceeding with the actual installation. See "About conditions for support of replicated NetBackup catalogs" on page 33. See "About replication considerations" on page 48.
Stage 2	Installing and configuring the clustered NetBackup primary on the main site	Install and configure the NetBackup primary server cluster on the main site. See "Installing and configuring the main NetBackup primary server cluster" on page 57.
Stage 3	Installing and configuring the clustered NetBackup primary on the alternate node	Install and configure the NetBackup primary server cluster on the alternate site. See "Installing and configuring an alternate NetBackup primary server cluster" on page 57.

About clustering considerations

The NetBackup primary server nodes on both the sites must be configured as clustered primary servers, although they can be single node clusters on each site.

The NetBackup primary server 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 main primary server requires a site failover operation.
Dual node on main site and single node on alternate site	This configuration requires three nodes—two at the main site and one at the alternate site. During normal operations, since there is single node at each site, site failover is not required to address issues with the main primary server. Instead a local failover can be used. However, there is no protection for the node on the alternate site. As a general best practice, this configuration is recommended.

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 the main NetBackup primary server cluster

To install the NetBackup primary server cluster follow the instructions that are described in the NetBackup Clustered Primary Server Administrator's Guide. Refer following guidelines to install the main NetBackup primary server cluster with catalog replication.

Installing the main NetBackup primary server cluster with catalog replication

1. During the installation of the NetBackup primary server cluster on the primary node, specify the following:
   - Replicated storage as the mount point for the cluster common storage.
   - All servers that are part of the domain.
   - The servers that form the alternate site cluster.
2. After the NetBackup cluster group is created, reconfigure the storage resources to include the replication control components.
3. For some replications layers, for example for Arctera Volume Replicator (VVR), the replication agent must be in a separate service group. You must link the agent with the NetBackup application service group.
4. If the replication technology includes a bandwidth planning and analysis tool, use this tool to assess the bandwidth requirements before you implement the replication layer. To estimate the replication traffic, install and configure the main primary server cluster and run backups for a few weeks. Use the analysis tool to measure I/O traffic and plan the replication layer based on the tool's recommendations before you implement replication to the alternate site.

Installing and configuring an alternate NetBackup primary server cluster

To install and configure the alternate NetBackup primary server cluster, it is not necessary to have the same number of nodes as the main NetBackup primary server cluster. The alternate NetBackup primary server cluster must be clustered, but can be a single node cluster. Refer following guidelines to install the alternate NetBackup primary server cluster with catalog replication.

To install an alternate NetBackup primary server cluster with catalog replication

1. Before starting the installation determine the mount point that will be used to mount the replicated volume from the main site.

   Do not mount the replicated volume at this time. Instead mount another formatted volume against this mount point. During the installation an empty catalog is created on this volume, which can later be discarded.

2. Install NetBackup on the alternate primary server cluster using the same virtual host name as the main primary server cluster. During the installation make sure to specify all media servers that are configured on the main primary server cluster as additional servers. This action ensures that the server list is consistent for both clusters.

3. After the installation is complete, shutdown NetBackup and take the cluster common catalog volume off line.

4. Update the cluster configuration to mount the replicated volume from the main site as the cluster common volume. This step may require adding the replication agent or enabling replication within the disk resource. Do not bring this resource online until you are instructed to do so in the next phase of the setup. The volume that is used during the installation is no longer required and can be re-used for another purpose.

Caution: Do not configure the replication agent on the alternate site to automatically reverse the direction of replication. Do not reverse the replication until the main site is operational again.

Populating the server tables in the NetBackup database

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

To populate the server tables in the NetBackup database

1. After the cluster setup on the alternate site is complete, take NetBackup offline on the main site.
2. Reverse the replication direction and bring NetBackup online on the alternate site.

   Bringing the NetBackup primary server cluster online on a particular node of the cluster automatically adds that node to the list of known servers in the server tables. If there are multiple nodes at the alternate site, you must fail over the primary server to each node.
3. After all the member nodes are added, take NetBackup offline from the alternate site.
4. Reverse the replication direction and bring NetBackup online on the main site.

Upgrading NetBackup in a clustered replication configuration

In order for global failover to work correctly, both the main and the alternate site clusters must run the same version of NetBackup. This means that both clusters must be upgraded at the same time. The upgrade process requires you to disable the replication link and upgrade each cluster independently.

To upgrade NetBackup in a clustered replication configuration

1. Disable global cluster failover for the duration of the upgrade.
2. Suspend replication between the main and the alternate sites.
3. Ensure that the alternate site cluster is isolated from the wider network so that it cannot contact the media servers and clients when it comes online.
4. Upgrade NetBackup on the main primary server cluster, following the standard upgrade procedures. (Refer to the NetBackup Upgrade Guide.)
5. If you have both NetBackup CA-signed and external CA-signed certificates, you need to manually configure the active and the inactive nodes after the DR package recovery. Certificates for the virtual name are backed up with the DR package during catalog backup.
6. Run backup and restore tests to confirm that the upgrade has been successful.
7. Bring the alternate site cluster online and mount the replicated catalog volume to it.
8. Upgrade NetBackup on the alternate primary server cluster following the standard upgrade procedures. (Refer to the NetBackup Upgrade Guide.)

9 Perform the following steps, based on the type of certificates in your environment.

NetBackup CA-signed (host ID-based) certificates	Fetch the host ID-based certificates for all nodes on the alternate site, following the procedure to deploy the certificate. (Refer to the NetBackup Security and Encryption Guide.)
External CA-signed certificates	For successful host communication, you must configure all nodes to use external certificates after a disaster. For more details, refer to the NetBackup Security and Encryption Guide.

10 Take the alternate site cluster offline. This action avoids unnecessary and potentially time-consuming post processing operations on the alternate site catalog volume.

11 Re-connect the alternate site cluster to the wider network.

12 Restart the replication process and allow the replicated volume to fully synchronize.

13 Enable global cluster failover.

14 Perform a failover to the alternate site and perform backup and restore tests to confirm that the upgrade was successful.

15 If required, fail the global cluster back to the main site.

Failing over to an alternate primary server cluster

You must failover to the alternate primary server cluster under the following circumstances:

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

The exact failover procedure can vary for different replication technologies. Refer the following procedure listing the high-level steps to failover to the alternate primary server cluster.

To fail over to an alternate primary server cluster

1 Stop NetBackup on the main primary server cluster.

2 Stop or reverse the replication of the catalog volume.

3. If necessary, update the DNS with the new virtual IP address for the primary server.
4. Start NetBackup on the alternate primary server cluster.

Note: Steps 1 and 2 occur automatically if the main site fails.

In globally clustered environment the process of failing over to alternate primary server can be automated. To automate the process, multiple heartbeat connections must exist between the clusters. Failure of the heartbeat network can cause the alternate primary server cluster to come online while the main primary server cluster is still operational.

Testing the NetBackup primary server cluster in clustered replication environment

It is recommended to test the ability of the alternate primary 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.

Following procedure provides the high-level steps that must be followed for testing.

To test the NetBackup primary server cluster in clustered replication environment

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

Chapter 6

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

This topic provides links to instructions on how to perform user-directed backups and restore data in a cluster. See the NetBackup Administrator Guides for information on performing backups and restores in non-cluster environments.

The backup and restore process is the same whether you are in a cluster or a non-cluster environment. See the 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.

About restoring data in a cluster

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 primary server. See the information on how to allow redirected restores in the NetBackup Administrator's Guide, Volume I.

There may be other situations that require the appropriate `altnames` directory entries to be created on the primary 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 64.

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 primary server.

   For a Linux server:

   ```
   /usr/openv/netbackup/db/altnames/tic
   /usr/openv/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 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 application agents 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. The NetBackup client is installed on a cluster as it is in a non-clustered environment. Refer to the *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.