



Implementing Highly Available Data Protection with Veritas NetBackup

Alex Davies | January 2008

Contents

1.0	INTRODUCTION	4
1.1	DISASTER RECOVERY AND HIGH AVAILABILITY – WHAT IS THE DIFFERENCE?	4
1.1.1	<i>High Availability</i>	4
1.1.2	<i>Disaster Recovery</i>	4
1.2	DISASTER RECOVERY – ALL OR NOTHING?	4
1.3	GLOSSARY	5
1.4	REFERENCES	5
2.0	DESIGN CONSIDERATIONS FOR EFFECTIVE DISASTER RECOVERY	7
2.1	SERVICE LEVEL AGREEMENTS AND HIGH AVAILABILITY	7
2.2	PLANNING FOR PROLONGED OUTAGES	8
2.3	TESTING THE DISASTER RECOVERY CAPABILITY	8
3.0	PROTECTING AGAINST COMPONENT FAILURES	9
3.1	INFRASTRUCTURE COMPONENTS	9
3.1.1	<i>Network Links</i>	9
3.1.2	<i>Storage Device Connections</i>	9
3.1.3	<i>Storage Devices</i>	10
3.1.4	<i>Media Availability</i>	10
3.2	SERVER COMPONENTS	11
3.2.1	<i>NetBackup Master Server</i>	11
3.2.2	<i>NetBackup Media Servers</i>	11
3.2.3	<i>BMR Boot Servers</i>	13
3.2.4	<i>NetBackup LAN Clients</i>	13
3.2.5	<i>NetBackup SAN Clients</i>	13
4.0	NETBACKUP CATALOG PROTECTION	14
4.1	TYPES OF CATALOG BACKUP	14
4.2	CATALOG PROTECTION – GOOD/BETTER/BEST	14
4.2.1	<i>Good – Scheduled Catalog Backup</i>	14
4.2.2	<i>Better – Catalog Backup under Vault Control</i>	15
4.2.3	<i>Best – Catalog Replication to the DR Server</i>	15
5.0	NETBACKUP CLUSTERING OPTIONS	16
5.1	NETBACKUP MASTER SERVER LOCAL CLUSTER CONFIGURATION	16
5.2	NETBACKUP MASTER SERVER GLOBAL CLUSTER CONFIGURATION	17
5.3	VIRTUAL STORAGE UNITS AND APPLICATION CLUSTERS	18
6.0	AVOIDING SINGLE POINTS OF FAILURE WITHIN BACKUP STORAGE	19
6.1	DUPLICATION OF BACKUP DATA UNDER NETBACKUP CONTROL	19
6.1.1	<i>Good – Duplication (Vault controlled)</i>	19
6.1.2	<i>Better – in-line copy (policy controlled)</i>	20
6.1.3	<i>Best – Storage Lifecycle Policies</i>	20
6.2	CONSIDERATIONS WHEN USING DISK STORAGE	20
6.3	DUPLICATION OF BACKUP DATA OUTSIDE OF NETBACKUP CONTROL	20
6.3.1	<i>Considerations for replication within VTLs</i>	21
6.3.2	<i>Considerations for replication within disk arrays</i>	21
6.4	ENCRYPTION OF DATA GOING OFF SITE	21
7.0	DISASTER RECOVERY – PROTECTING AGAINST SITE FAILURE	23
7.1	COPY SELECTION FOR RESTORE	23

White Paper: Implementing Highly Available Data Protection with NetBackup

- 7.2 GOOD – DEDICATED DISASTER RECOVERY SITE WITH VAULT..... 23
 - 7.2.1 *Benefits of a dedicated disaster recovery site*..... 24
 - 7.2.2 *Limitations of a dedicated disaster recovery site* 24
- 7.3 BETTER (A) – RECOVERY WITHOUT IMPORT 25
 - 7.3.1 *Benefits of recovery without import* 26
 - 7.3.2 *Limitations of recovery without import*..... 27
- 7.4 BETTER (B) – MEDIA SERVER PROMOTION 27
 - 7.4.1 *Benefits of Media Server promotion* 28
 - 7.4.2 *Limitations of Media Server promotion*..... 28
- 7.5 BETTER (C) – GLOBALLY CLUSTERED REPLICATED MASTER SERVER 29
 - 7.5.1 *Benefits of a replicated Master Server cluster* 30
 - 7.5.2 *Limitations of a replicated Master Server cluster* 30
- 7.6 BEST – DUAL SITE/SINGLE DOMAIN 30
 - 7.6.1 *Benefits of the Dual Site/Single Domain model* 32
 - 7.6.2 *Limitations of the Dual Site/Single Domain model* 32
- 7.7 DISASTER RECOVERY OPTIONS FEATURE COMPARISON 33

1.0 Introduction

The data protection system must be regarded as a 'mission critical' element of any modern data center. The design of any data protection system must, as far as possible, eliminate single points of failure so that data can be recovered to an acceptable state and point in time in the event of data, server or site loss.

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

Although this paper has been written for NetBackup 6.5 GA and NetBackup 6.5.1 many of the concepts described here can also be applied to NetBackup 6.0. It uses a "good"/"better"/"best" classification in certain areas to indicate the merits of one solution over another. In general these classifications may be regarded as follows:

Good – provides an adequate solution in smaller environments and environments where data is less critical.

Better – provides a more robust solution, generally at increased cost or complexity.

Best – provides the best currently available solution but often at significant cost.

1.1 Disaster Recovery and High Availability – what is the difference?

This paper discusses the topics of high availability (HA) and disaster recovery (DR). It is important to understand that these are not the same. While high availability technologies may form part of a disaster recovery solution, simply having high availability components does not ensure that disaster recovery is possible

1.1.1 High Availability

High availability solutions exist to protect against single points of failure within an environment, thus ensuring that the environment continues to operate as intended (although possibly at reduced capacity) in the event of the failure of a component. In a NetBackup domain high availability can take many forms ranging from the clustering of a Master Server to the use of the Shared Storage Option and NIC teaming to protect against tape drive and network failures.

1.1.2 Disaster Recovery

Disaster recovery is the general term used to describe a process involved in the restoration of normal (although possibly reduced) service following a major unforeseen event. In the context of this paper, disaster recovery means the process involved in restoring data backed up by NetBackup to servers in a different location following the loss of a data center or site.

1.2 Disaster Recovery – all or nothing?

The traditional view of disaster recovery is that the complete environment must be re-created at a secondary site but this is often prohibitively expensive. In practice the immediate priority in the wake of a disaster involves the recovery of key 'mission critical' systems only and does not require a complete recovery of the entire environment. Organizations are increasingly seeking ways of providing a disaster recovery capability without needing to have a complete mirror of the production environment sitting idle in case a disaster occurs. This paper looks at a variety of approaches to meeting the objective of providing rapid recovery capability without incurring the expense of a fully dedicated disaster recovery site or the time penalty of building a recovery environment from the ground up.

1.3 Glossary

The following terms are used throughout this document:

Data Protection Solution – While a data protection solution may encompass many different technologies and levels of protection, in the context of this document it means backup and recovery provided by NetBackup.

Domain – A NetBackup Domain is a collection of NetBackup Servers and Clients under the control of a single Master Server

Site – A site is a single data center, data hall, building or location where servers requiring backup are located. A single domain may cover two or more sites. Three sub-classes of sites are described in this document:

Primary site – this is the site at which the servers that are protected by the data protection solution normally reside and is also described as the 'production' site.

Secondary site – this is the site at which data from the primary site would be recovered in the event of a disaster. The secondary site may be a dedicated disaster recovery site or a second production site.

Disaster recovery site/domain – this is a dedicated facility at the secondary site that is used to recover the data from the primary site in the event of a site loss. This may be either a separate NetBackup Domain or part of another existing NetBackup Domain.

NetBackup catalog – NetBackup catalogs are the internal databases that contain information about NetBackup backups and configuration. Backup information includes records of the files that have been backed up and the media on which the files are stored. The catalogs also contain information about the media and the storage devices. In the event of a disaster in which the NetBackup Master Server is lost a full or partial recovery of the catalog information is required before recovery from backup can begin.

Single Point of Failure (SPOF) – any component of the data protection solution which, if it fails or is unavailable, prevents the backup and recovery process from working.

SPAS – Symantec Product Authentication Service, formerly known as VxSS.

RPO (recovery point objective) – the most recent point an application can be recovered to using the available backup data. The actual recovery point may differ from any established objective. Proper planning and testing may need to be executed in order to ensure that the actual recovery point aligns with the desired recovery point objective.

RTO (recovery time objective) – the time required to recover an application, server or series of applications or servers following a failure. In the context of this document, the RTO is generally assumed to be the RTO at the secondary site, including the time to prepare the NetBackup environment at the secondary site to begin restoring data. Again, it should be noted that the actual recovery time may differ from any established objective. Proper planning and testing may need to be executed in order to assure that the actual recovery time aligns with the desired recovery time objective.

1.4 References

Many of the concepts discussed in this paper are discussed in more detail in other documents, in particular:

1. Information about how to cluster NetBackup Master Servers can be found in the "NetBackup High Availability Administrators Guide".
2. Information about the various commands described in this paper can be found in the "NetBackup Commands for Windows" and "NetBackup Commands for UNIX and Linux" guides.

White Paper: Implementing Highly Available Data Protection with NetBackup

3. Information about how to restore catalog backups can be found in the “NetBackup Troubleshooting Guide”.
4. Information on SPAS, NBAC, Client Encryption and Media Server Encryption can be found in the “Security and Encryption Guide”.
5. Details of the support requirement from globally clustered Master Servers can be found in tech note 287636.

All of these documents, together with the NetBackup compatibility lists, can be found at the NetBackup support web site at <http://www.symantec.com/business/support>.

2.0 Design Considerations for effective Disaster Recovery

When designing data protection systems some aspects are routinely overlooked, resulting in a solution that fails to deliver the intended response times or degrees of recoverability. This section looks at three of these aspects.

2.1 Service level agreements and high availability

When setting up a data protection system it is important to ensure that the recovery service levels are clearly defined and achievable to avoid any misunderstanding or misinterpreting the recovery expectations on the part of the end users (the 'customers' of the data protection system). It is surprising how often assumptions are made about recoverability that are simply unachievable with the given data protection infrastructure. These assumptions are often made by the people operating the system themselves and not just the end users.

One common mistake in this area is to believe that making a backup every day guarantees a 24 hour recovery point. Clearly if an application is backed up on a daily basis and the backup takes 2 hours to complete the time between the start of yesterday's backup (the last valid recovery point) and the end of today's backup is at least 26 hours. This stretching of the RPO is tolerable provided the backup runs successfully every night, but assume for a moment that the backup starts at 8:00 PM and there is no overnight operational coverage of the data protection system. The backup may fail at 9:00 PM and the failure may not be discovered until 9:00 AM the next day. Even if the backup is then successfully re-run the RPO has extended from a tolerable 26 hours to a potentially unacceptable 39 hours. If the backup cannot be re-run immediately, for example due to hardware failure, then this period is stretched even further, exposing the business to ever increasing exposure to data loss.

It should be apparent that even in the most highly available data protection environment a recovery point service level of 24 hours cannot be guaranteed if only one backup is made each day. However, the deployment of a data protection solution made up of suitably resilient components can ensure a high rate of backup success and increase the probability that an acceptable RPO can be maintained.

Another common mistake is a failure to set a realistic recovery time expectation in the service level agreement. The total RTO figure should include allowances for the following factors:

- Time to detect a problem has occurred and decide on the recovery approach
- Time to provision new servers and storage
- Time to recover the backup/restore infrastructure
- Time to restore the backup (including both full and incremental backups up to the last backup)
- Time to roll the application forward to the RPO

In particular it is important to bear in mind that the time required to restore an environment from a series of full and incremental backups will be greater than the time required to recover from a single full backup.

As the factors listed above will vary depending on the nature of the failure scenario it is advisable to have several different agreed levels of RTO to cover different eventualities. For example, the RTO for an application or server failure may be shorter than the RTO for a site failure.

Finally, both RTO and RPO values should be realistically achievable. From an operational viewpoint these should be 'worst case' rather than 'best case' figures (i.e. the expectation should be that, in most cases, these values will be exceeded). As illustrated above, an RPO of 24 hours is not realistic if the application is backed up on a daily basis – particularly if there is no

real time monitoring for failures. Likewise an RTO of 24 hours is not realistic if it takes 48 hours to acquire and provision a replacement server.

2.2 Planning for prolonged outages

A full scale disaster such as a site loss is likely to render the primary or production data center inoperable for a significant period of time. When designing the disaster recovery model it is important to remember that it will be necessary to continue making backups at the disaster recovery site and, ultimately, be able to restore those backups at the production site. Consideration needs to be given to things like the storage capacity at the disaster recovery site (how big is the tape library, etc.), as well as how easily reversible the fail over process is.

2.3 Testing the disaster recovery capability

It is important to periodically test the disaster recovery process to ensure that it works in practice as well as in theory. When designing the data protection solution it is important to consider how the disaster recovery plan will be tested and how often testing will occur. The more frequently and comprehensively the plan is tested the more successful its execution in a real disaster is likely to be. However, comprehensive testing is not required if the process is well documented and understood. The various site disaster recovery approaches described in a later section of this paper have different test requirements and complexities which may influence the choice of approach used.

3.0 Protecting against Component Failures

A data protection system consists of a number of different components, each of which has the potential to fail and disrupt the backup or restore process. This section looks at these various components and how single points of failure can be eliminated at each component level either by making the component itself highly available or by deploying multiple components for redundancy.

3.1 Infrastructure components

3.1.1 Network Links

The majority of backup traffic is transferred over network connections with 100 Mbit and 1 Gbit networks (typically giving 8 Mbytes/sec and 65 Mbytes/sec respectively). Network links can be made highly available by deploying redundant network teaming. The cost of such teaming means that it 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) are accepted.

3.1.2 Storage Device Connections

Connections to storage devices and their controllers also represent single points of failure. If the connection is not present the device cannot be used.

3.1.2.1 SAN connected tape and disk devices

SAN connections generally exist between the backup servers and the backup storage, although the NetBackup SAN Client also supports the concept of SAN connections from clients to Media Servers. In all cases, SANs should always be configured to provide redundant connections between the source and target components.

Most SAN attached disk arrays have redundant SAN connections and support dynamic multi-pathing (DMP) software that ensures 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, appear to servers as two separate devices. Prior to the release of NetBackup 6.0, if both device paths were presented it was necessary to mark one path as unavailable to NetBackup to prevent device allocation conflicts. Multi-path awareness was introduced in NetBackup 6.0 and NetBackup now recognizes that both device paths map to the same physical device. 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.

3.1.2.2 Robotic control connections

In tape based backup environments the robotic control also represents a single point of failure as the inability to send instructions to the tape library will prevent backup and restore operations, even if the tape drives are available.

Some types of tape libraries use dedicated control software that runs on a server independent of NetBackup (for example Sun STK ACSLS or Quantum's ATM) which may be clustered. The NetBackup Media Servers send requests to this server which, in turn, handles the movement to tapes between slots and drives in the tape library.

Other types of libraries depend on a direct device connection from one of the NetBackup servers (usually the Master Server) which provides the control instructions to the library. If this device connection is lost, the library cannot be used. Like tape drives, many SAN attached tape libraries support multiple connections to the robotic control for redundancy (some SCSI attached libraries also offer this capability). It is possible to configure these connections to provide

protection against server failure, (for example, by configuring one path to each node of a clustered Master Server), but care must be taken to ensure that both paths are not active at the same time as this could result in conflicting instructions being issued which could result in backup failure and data loss.

NetBackup does not currently recognize multi-path connections to tape libraries and path configurations must be switched manually in the event of a failure. Support for multi-path connections to tape libraries is expected to be included in the NetBackup 6.5.2 release.

3.1.3 Storage Devices

Storage devices, whether they are tape drives or disks, can and do fail. The key requirement here is to have multiple devices as backup targets.

A Media Server that has access to only one tape drive cannot complete backups to tape if that tape drive is down – Media Servers should be configured to access at least two tape drives. Using SAN attached tape drives which can be shared between Media Servers is an effective way of ensuring that tape drives are accessible without needing to provide large numbers of redundant devices. Typically one or two ‘spare’ drives may be provided for resilience and to allow restore operations while backups are in progress. For example, if four Media Servers share five tape drives and are each configured to use one drive then each Media Server can access one drive, even if one of the five is defective. In practice, it is more likely that each Media Server will be configured to use multiple tape drives so that all five drives are used as much as possible and the loss of a single drive will simply mean that the backup take longer to complete. If the Media Servers run backups at different times then the ratio of tape drives to Media Servers may be even lower without risking backup failure (although the loss of a drive will slow the backup process down).

In a similar way, SharedDisk Disk Pools can be created in NetBackup 6.5, allowing several Media Servers to share a common pool of disk storage, and AdvancedDisk Disk Pools can be created on individual Media Servers to protect against the failure of a single disk device.

3.1.4 Media Availability

In tape based backup solutions, failures can occur if no suitable tape media is available for use by a backup job. NetBackup incorporates two features to reduce the risk of failures of this kind, global scratch pools and media sharing.

3.1.4.1 Global Scratch Pools

All backup and duplication jobs writing to tape use tapes that are in a specific media pool and have the same retention criteria as the data being backed up. If no suitable tapes are available the backup will fail.

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. When a backup or duplication job runs and no suitable tapes are available in the media pool specified by the job an unassigned tape is transferred from the global scratch pool to the specified media pool and 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 the job specifies.

3.1.4.2 Media Sharing

Prior to the introduction of media sharing in NetBackup 6.5, tapes were assigned to one Media Server while they contained valid backups and could only be used by a different Media Server when all the backups had expired and the tape had been de-assigned.

The concept of media sharing allows partially full tapes to be used by different Media Servers until they are full. Only one Media Server can write to the tape at one time but, when the tape is

not in use it can be used by any Media Server requiring a tape from the particular media pool and with the same retention characteristics.

To complement the media sharing feature a value of 'maximum partially full media' has been introduced for media pools. Setting this value enforces media sharing by restricting the number of partially full tapes in a given media pool. Until one tape is full, another empty tape cannot be assigned to the pool. One effect of this feature is to ensure that tapes are used as efficiently as possible, however, another effect is to delay the start of backups if tapes are in use by other Media Servers.

3.2 Server components

3.2.1 NetBackup Master Server

There is only one Master Server for each NetBackup domain which controls all the backup activity within the domain. As such, the Master Server represents the most obvious single point of failure in the data protection environment – without it backups and restores are not possible.

For this reason the Master Server should always be made highly available in situations where the loss of the data protection system for a period of more than a few minutes is considered an unacceptable risk. NetBackup supports the clustering of the Master Server with a wide range of clustering tools including Veritas Cluster Server, Microsoft MSCS, SunCluster, HP MC ServiceGuard and IBM AIX HACMP. The clustering process is described in more detail in section 5.0.

3.2.2 NetBackup Media Servers

Although Media Servers may be configured with redundant network and SAN connections, the servers themselves remain single points of failure. Media Servers can be categorized into three distinct groups:

- Dedicated Media Servers which only run the Media Server software and exclusively back up data from other systems.
- Non-dedicated Media Servers which also run other applications that require backing up but also backup data from other systems.
- SAN Media Servers which also run other applications that require backing up but do not backup data from other systems.

3.2.2.1 Dedicated Media Servers and Storage Unit Groups

While it is possible to cluster Media Servers this is not generally necessary for dedicated Media Servers. Clustered configurations tend to be active/passive and require additional software components. Instead, Media Server redundancy can be achieved by simply allowing backups to run to more than one Media Server. This gives the additional benefit of an active/active operation.

In NetBackup, Storage Unit Groups can be used to provide not only protection against the failure of a single Media Server but also load balancing across multiple Media Servers to ensure optimal backup and restore performance. Configuration options allow Storage Unit Groups to behave in four different ways:

Failover – this mode always uses the first Storage Unit listed unless the Media Server is down. Excess jobs queue rather than being directed to the next Storage Unit. This behaviour is similar to what would be seen if two Media Servers were configured as an active/passive cluster.

Prioritized (good) – this mode uses the first available Storage Unit in the list. In this case jobs that exceed the total number that the Storage Unit can handle will be directed to the

White Paper: Implementing Highly Available Data Protection with NetBackup

next one in the list, and if the Media Server is down, all backups will be directed to the next one.

Round Robin (better) – this mode uses a different Storage Unit for each job and cycles around the list of Storage Units. If each Storage Unit is on a different Media Server, this acts as a simple load balancing mechanism.

Load balanced (best) – this mode is introduced in NetBackup 6.5 and only works with Flexible Disk and Media Manager Storage Unit types. In this mode, NetBackup carries out checks on activity and resources available on each Media Server before directing the backup to the one with the lightest load.

One best practice tip when using 'prioritized' and 'failover' groups is to configure two Storage Unit Groups to use two Media Servers as follows:

- 1) Each Media Server is configured to have a single storage unit so NodeA has STUA and NodeB has STUB.
- 2) Two storage unit groups are configured with the storage units in a specific order in each one, so SUGAB contains STUA followed by STUB and SUGBA contains STUB followed by STUA
- 3) Backup policies are then evenly shared between SUGAB and SUGBA

During operation, the backup traffic will normally be shared between the two nodes but if one node should fail, all backups will automatically go to the other node.

3.2.2.2 Non-dedicated Media Servers

Storage Unit Groups may also be used with non-dedicated Media Servers. However, there may be a requirement to provide high availability for the other applications that run on the Media Servers, resulting in clusters of two or more Media Servers. In this case, the Media Server software is still installed locally on each member node of the cluster and the servers defined as members of an 'application cluster' in the NetBackup EMM database using the name of the cluster as the application cluster name (see section 5.3). A storage unit may then be created using the virtual name of the cluster as the Media Server and the application can be backed up using this storage unit – which will always map to the active node of the cluster.

3.2.2.3 NetBackup SAN Media Servers

Unlike regular Media Servers, SAN Media Servers are only intended to protect themselves. A SAN Media Server connects directly to the backup storage in the same way as a regular Media Server but does not receive data from other client systems over a network or SAN link.

SAN Media Servers are usually deployed on servers supporting large, mission critical applications, which are often clustered. While the application may be clustered, it is not necessary to cluster the SAN Media Server itself. Instead the SAN Media Server software should be installed on each member node of the cluster and the servers defined as members of an 'application cluster' in the NetBackup EMM database using the name of the cluster as the application cluster name (see section 5.3). A storage unit may then be created using the virtual name of the cluster as the Media Server and the application can be backed up using this storage unit – which will always map to the active node of the cluster.

3.2.2.4 Restoring backups using an alternative Media Server

In most cases when a restore is made NetBackup expects the same Media Server and client to be used for the restore operation that was used for the original backup. In a disaster recovery situation, it is often necessary to use a different Media Server to restore the backup to a different client as the Media Servers and clients at the disaster recovery site are likely to have different names to those at the primary site.

NetBackup allows the configuration of 'failover restore Media Servers' that can be used to handle restores in the event the original Media Server is down. These are configured under the host properties -> Master Server -> Restore Failover tab in the administration GUI or, in the case

of UNIX and Linux Master Servers, by creating `FAILOVER_RESTORE_MEDIA_SERVER` entries in the `bp.conf` file.

3.2.3 BMR Boot Servers

If bare metal restore (BMR) is being used for server recovery (a best practice recommendation), then multiple boot servers should be configured within the domain to ensure that recovery is still possible when one server is unavailable. For site disaster recovery, BMR boot servers must also exist at the disaster recovery site. These requirements can be avoided if CD based SRTs are used instead of boot servers.

3.2.4 NetBackup LAN Clients

Where network attached, or LAN clients are involved there is no need to cluster the NetBackup client software. Simply installing the client software and application agent software (if required) on each node of the cluster and then backing the application up using the virtual server name should be sufficient to ensure a good backup is made on the active node of the cluster.

3.2.5 NetBackup SAN Clients

The NetBackup SAN Client is a new feature introduced in NetBackup 6.5. Like the SAN Media Server, it avoids sending backup traffic over the network to the Media Server, but instead of sending the data directly to backup storage; it sends it across a SAN link to a Media Server. The primary advantage of using a SAN Client over a SAN Media Server is that most of the backup processing is transferred to a remote Media Server so the impact on the running application is reduced.

Like SAN Media Servers, SAN Clients are often used to protect clustered applications. When used in this way they should also be configured as 'application clusters' in EMM (see section 5.3). Doing this 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.

4.0 NetBackup Catalog Protection

The NetBackup databases that reside on the Master Server (collectively known as the 'NetBackup catalog') contain all the information about the contents of all the valid backups in the domain. Without this information it is impossible to restore anything. The key elements that make up the catalog in NetBackup 6.x are:

The image database – this records the contents of all backups and the storage devices they are written to and is structured by a client. This is the database that is browsed during restore operations.

The policy or class database – this contains information about the backup policies including the schedules, file lists, clients and storage destinations related to each policy

The EMM database – this database was introduced in NetBackup 6.0 and replaces a number of separate databases that existed in earlier versions of NetBackup. It contains information about devices, storage destinations, media, servers and various other infrastructure components.

The BMR database – this database contains the BMR data collected from each client when a backup policy is run and used during a BMR recovery to automatically restore the backups associated with that client.

While clustering can be used to make the Master Server highly available this does not protect the catalog data itself and it is important to ensure that it is backed up on a regular basis or replicated to protect against site loss and storage failure.

4.1 Types of Catalog Backup

NetBackup 6.x offers two types of catalog backup, depending on the version of NetBackup used:

Online or hot catalog backup – This type of catalog backup was introduced in NetBackup 6.0. It is considered an 'online' or 'hot' method because it can be performed while regular backup activity occurs. This is policy-based backup job and can span more than one tape. It also allows incremental backups of the catalog, which can significantly reduce catalog backup times for large catalogs.

Offline or cold catalog backup – This is the traditional type of catalog backup available for all versions of NetBackup up to and including 6.x. It is considered an 'offline' or 'cold' backup because running a backup of this type takes NetBackup 'offline' preventing backup and restore activity from starting while the backup is running. (In NetBackup 6.x the Sybase ASA databases (NBDB and BMRDB) are shut down during the backup.) This type of catalog backup must fit on a single tape and alternates between two specified tapes or disk locations so only the two most recent catalog backups are retained. It also does not allow incremental backups. This type of catalog backup is now considered obsolete and will be discontinued after NetBackup 6.x.

Note: Catalog backups may only be written to tape or BasicDisk storage and cannot be written to OpenStorage appliances, PureDisk Disk Option or Flexible Disk Option storage.

4.2 Catalog Protection – Good/Better/Best

The good/better/best model can be applied to protecting the catalog as described in the following sections.

4.2.1 Good – Scheduled Catalog Backup

Both the online and offline catalog backups can be scheduled to run at regular intervals.

Scheduling the online catalog backup – the online catalog backup is scheduled using a NetBackup catalog backup policy. This is defined like any other NetBackup policy with a regular schedule and storage unit but is subject to the following limitations:

White Paper: Implementing Highly Available Data Protection with NetBackup

1. The backup must run to tape or a BasicDisk Storage Unit
2. If run to tape, the media pool used must be identified as a 'catalog backup' pool.

The time that the backup runs should ideally be scheduled to capture the state of the catalog at the end of the daily backup run. However, as the running of the catalog backup does not prevent other backups running, it may be linked to other events such as the daily shipment of tapes to off-site storage.

One useful feature of the online catalog backup policy is that it can be used to record and e-mail details of the tapes used by policies that back up critical applications so that the presence of these tapes at the disaster recovery site can be confirmed before restore operations are started.

Scheduling the offline catalog backup – the offline catalog backup is an event triggered operation and can be set to occur on three different triggers:

1. After each session of scheduled backups – the catalog backup is automatically initiated shortly after the scheduled job work list is completed. In most cases this means that a catalog backup occurs at the end of the daily backup run and captures information about the latest backups created.
2. After each session of scheduled, manual and user backups – in addition to running at the end of the daily backup run, a catalog backup is also triggered after ad-hoc backups are run. This results in more backups but ensures that the information is more up to date.
3. Manual initiation – this relies on a backup administrator or external script starting the backup. This option is sometimes used to decrease the frequency of backups to ensure that the catalog backup does not prevent scheduled backups from executing. For example, when the daily backup run is split into several blocks or periods and only one daily catalog backup is required.

4.2.2 Better – Catalog Backup under Vault Control

The Vault Option includes the ability to run a catalog backup job before the duplicated media are ejected from the tape library – this can be done whether or not the Vault Option is actually used to manage the duplication process. Placing the catalog backup under the control of the Vault Option ensures that the catalog backup is created that includes information about the duplicated tape copies and sent off-site with them. The catalog backup used to recover the Master Server at the disaster recovery site will include all the information about the duplicate tapes, simplifying the recovery process.

Note: The Vault Option only supports a hot catalog backup as part of the vaulting process.

4.2.3 Best – Catalog Replication to the DR Server

Replicating the catalog to a standby Master Server at the DR or secondary site offers the best solution for rapid catalog recovery at the disaster recovery site. Continuous replication ensures that the catalog is as up to date as the replication link allows. Replication of catalogs is discussed in more detail in section 7.0. However it should be noted that replication does not protect against catalog corruption or the effects of accidentally deleting or expiring images and regular scheduled catalog backups should also be made to protect against this. Thus the true 'best' configuration is a combination of replication and scheduled catalog backup.

5.0 NetBackup Clustering Options

This section looks at the three different types of clustering that are discussed elsewhere in this document. Clustering to control the NetBackup software is normally only applied to NetBackup Master Servers but non-dedicated Media Servers, SAN Media Servers and SAN Clients may also make use of application clusters within NetBackup's EMM database to relate the physical nodes running the NetBackup software to a virtual node supporting a clustered application.

5.1 NetBackup Master Server Local Cluster Configuration

The Master Server represents a single point of failure within a NetBackup domain and it is generally recommended that it be configured to be highly available using a clustering solution. NetBackup supports the use of a wide range of clustering technologies to cluster the Master Server including Veritas Cluster Server (VCS), Microsoft Cluster Server (MSCS), MC ServiceGuard, SunCluster and HACMP, that allow clustering on all platforms that support a NetBackup Master Server (please refer to the cluster compatibility HCL for details of specific versions of NetBackup, clustering software and operating system supported).

Clustering of a Master Server is done at installation time using tools provided by NetBackup. For UNIX/Linux platforms the installation process to cluster a Master Server involves the following steps:

1. Preparing the cluster environment (often called seeding the cluster) by installing the clustering software on the appropriate nodes and establishing basic cluster functionality.
2. Installing the NetBackup Master Server software on each node of the cluster specifying the virtual server name as the Master Server name (this ensures that the NetBackup clustering agent software is installed).
3. On one node of the cluster executing `cluster_config` script located in `/usr/opensv/netbackup/bin/cluster`. This script guides the user through the cluster configuration process.
4. In some cases adjusting the cluster configuration such as adding additional network interfaces, may be carried out after the cluster has been configured.

For Windows (using VCS or MSCS), the installation process is slightly different. The initial seeding of the cluster is still done in the same way but the NetBackup installation wizard detects the presence of the cluster and pushes the software out to all the member nodes specified during the installation.

Note: Clustering must be done at the time of the installation and a standalone Master Server cannot be clustered retrospectively without the assistance of Symantec Professional Services.

More details about clustering NetBackup Master Servers can be found in the NetBackup High Availability Guide.

Advances in clustering and disk mirroring technologies mean that it is now possible to use a local VCS cluster to link sites up to 20 miles (30 km) apart. In this configuration any replication layer is fully synchronous and hidden from the storage presentation which appears to be a full mirror. As the cluster is local rather than global, all nodes are required to be on the same layer 3 networks and will require multiple dedicated heartbeat connections. These requirements often prevent local clusters from being configured over such distances and they are usually limited to a single data center or separate 'data halls' in the same facility.

Note: As the replication layer is fully synchronous, performance decreases with separation – 20 miles of separation is the maximum practical limit for such a configuration and a Master Server cluster with storage mirrored over this distance will run significantly slower than one in which the storage is mirrored within the same array.

5.2 NetBackup Master Server Global Cluster Configuration

A global cluster is constructed of two or more individual local clusters connected together using the Veritas Cluster Server Global Cluster Option (GCO) to form a single cluster. Each local cluster may consist of one or more nodes (where one node is used the cluster configuration is known as a 'single node cluster').




Figure 1 – Global Cluster

Figure 1 above shows a three node global cluster with two nodes at the primary site and one at the secondary site. The good/better/best model can be applied to the number of nodes in a global cluster as follows:

Good – two nodes configured as single node clusters, one at each site. This is the bare minimum configuration required and assumes that the cluster will always be failed over to the secondary site, irrespective of the nature of the problem affecting the primary site.

Better – three nodes configured to provide a two node active/passive local cluster at the primary site with a third node, configured as a single node cluster, at the secondary site. This configuration allows local failover at the primary site in the event of individual server failure or for routine maintenance purposes as well as global failover to the secondary site in the event of site or storage loss at the primary site. However, once failed over to the secondary site there is no protection against server failure as there is no standby node at that site. This configuration generally offers the best cost/resilience compromise.

Best – four nodes configured to provide two active/passive local clusters, one at each site. This configuration provides protection against server failure at both the primary and secondary sites without needing to carry out a site failure. It also ensures protection against a single point of failure at the secondary site in the event that the primary site has been lost.

In global cluster configurations the storage containing the NetBackup catalogs is replicated between the sites rather than mirrored as it is in local clusters (although mirroring of the storage within the arrays at both the primary and secondary sites should still take place). The replication technology used may be either hardware (array/SAN) or software (server/network) based but must be continuous and preserve write order fidelity – snapshot technologies are not supported for replication of the storage in this model. Although the replication must be continuous, it does not need to be fully synchronous and asynchronous replication that preserves write order fidelity allowing separation over greater distances than the local clustering model allows. Depending on the replication technology used and replication latency separation distances of hundreds of miles are possible.

5.3 Virtual Storage Units and Application Clusters

The NetBackup EMM database, introduced in NetBackup 6.0, includes information about the various servers known to NetBackup including whether or not the nodes form part of a cluster.

Application Cluster definitions in EMM are used to denote the fact that nodes are members of cluster. They don't provide clustering capability directly, but do ensure that connections to a cluster are always directed to the active node of the cluster.

In a typical NetBackup environment, the storage unit is associated with a NetBackup entity, such as the Master Server or Media Server. However, if Media Servers or SAN Media Servers are installed on the member nodes of a cluster and an application cluster is created within NetBackup using the same virtual name as the clustered application, a storage unit may also be associated with the virtual name of the application in the cluster environment.

The storage unit for that application is not associated with any specific node in the cluster and can failover to any node in the cluster. This guarantees that the storage unit follows along with the application, wherever the application is running in the cluster. The NetBackup Master Server uses the virtual name of the application as the Media Server name when it backs up the clustered application.

Application clusters can also be used with SAN Clients in the same way to ensure that the fiber transport connection from the Media Server always goes to the node of the cluster where the clustered application is running.

An Application Cluster consists of a host record of type `app_cluster` in the EMM database with which existing physical host records are associated and is defined using the `nbemmcmd` command (located in `<install path>/bin/admincmd`) as follows:

To add the cluster name to the EMM database:

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

The name `<cluster>` here should be the virtual host name associated with the clustered application to be backed up (where two or more virtual names exist it will be necessary to create multiple application clusters).

To add the member nodes to the cluster (assuming records for them already exist in EMM) use the following command:

```
nbemmcmd -updatehost -add_server_to_app_cluster -machinename <member> -  
clustername <cluster> -netbackupversion <version number> -masterserver <Master  
Server>
```

The member nodes should also have their NetBackup Server list updated to include the virtual host names as well as the physical ones.

An application cluster should always be configured when non-dedicated Media Servers, SAN Media Servers or SAN Clients are used to protect a clustered application.

6.0 Avoiding Single Points of Failure within Backup Storage

Even if there are no single points of failure within the backup and restore process the location of the backed up data may constitute an exposure as the inability to gain access to the storage or to be able to read from it prevents the data from being restored.

Best practice dictates that all data should exist in at least two places. This is why we create backups in the first place but, once created, the backed up data needs to be transferred to a remote location to ensure recoverability in the event of a site loss. Keeping backups in a tape library that sits along side the disk array containing the application data (or even to another disk within the same array) is no use if the building or room that they sit in is destroyed or rendered inaccessible.

However, transferring the backups to another location may significantly extend the recovery time in the event of a less catastrophic failure such as a data corruption, disk failure or accidental deletion of data. For this reason it is advisable to create multiple copies of the backup data and store them in different locations.

6.1 Duplication of Backup Data under NetBackup Control

When duplication is carried out under the control of NetBackup, details of the location of each backup copy (i.e. which tapes or disks the backup is held on) are automatically recorded in the NetBackup catalog ensuring that NetBackup is 'aware' of all existing copies.

Three methods of duplication under NetBackup control are described in this section and it is important to understand the difference between them:

In-line copy – this involves writing two or more copies of the backup at the same time and is controlled by the number of copies set in the original backup policy definition.

Duplication – this involves copying the backup from one Storage Location to another once the initial backup has completed using the 'bpduplicate' function within NetBackup.

Optimized duplication – some OpenStorage appliances offer replication capabilities and allow NetBackup to take advantage of a feature known as 'optimized duplication'. This is a variant of duplication in which replication between storage destinations takes place at the appliance level outside of NetBackup control. When a NetBackup duplication instruction is issued, the appliance checks the status of the replication operation, and, if it has completed, the NetBackup catalog is updated to record the presence of a second copy at a second location.

The advantages and disadvantages of each method are described in the following sections. Using these two methods, NetBackup offers three distinct mechanisms for duplication which can be aligned to a good/better/best model as shown here:

6.1.1 Good – Duplication (Vault controlled)

The NetBackup Vault Option can be used to control duplication of backups. Duplicating after the backup has completed means that only one backup device is required to create the initial copy of each backup and one advantage often cited for duplication is that it 'sweats the assets' by using tape drives at times when they might otherwise be idle, thus reducing the amount of time expensive equipment sits idle. While this is true, the down side of duplication is that there is a substantial lag between the creation of the original backup and the creation of a copy that can be sent off-site. During this time the exposure to site loss remains present.

6.1.2 Better – in-line copy (policy controlled)

The advantage of in-line copy over duplication is that there are always two copies of the backup so a copy can be sent to off-site storage as soon as the backup completes (or more typically as soon as the backup run completes). The main disadvantages of in-line copy are that the same Media Server must be used to create both copies and, in certain situations, the overall performance of the backup can be degraded by the process (this only occurs in situations where the read speed of the source data being backed up is relatively fast). A further disadvantage when compared to duplication is that more storage devices are required to complete the backups within a designated 'backup window' – in an all tape environment twice as many tape drives are required, which may represent a significant capital outlay.

In-line copy can be configured so that if one copy fails the backup still runs to completion, but if this option is taken, the additional copy must be created manually as a duplication process which can be both complex and time consuming.

6.1.2.1 In-line copy performance

When in-line copy is used data is read once into the Media Server and then written out twice, once to each storage device. This can result in a backup appearing to run slower using in-line copy than writing to a single copy if the source data can be read faster than the two targets can be written to. This does not normally present a problem but in some cases it may have the effect of extending the initial backup time. However, as there is no need for subsequent duplication of the backup, the total time to create all copies of the backup is still less than it would be if duplication is used.

6.1.3 Best – Storage Lifecycle Policies

Storage Lifecycle Policies, introduced in NetBackup 6.5, provide an effective means of automatically ensuring that backup data is written and duplicated to multiple storage locations, thus avoiding the potential risk of losing the only copy of the backup that exists. Storage Lifecycle Policies can be configured to do both in-line copy and duplication or a mixture of both.

Unlike Vault based or manually scheduled duplication, which is usually run outside of the backup window, duplication initiated by Storage Lifecycle Policies usually starts at the first practical opportunity after the backup completes, ensuring a shorter period of exposure when only a single copy of the data exists.

A key advantage of Storage Lifecycle Policies over in-line copy is that if an in-line copy job fails the Storage Lifecycle Policy will automatically create a duplication job to ensure that the copy is made. As this duplication job starts almost as soon as the original backup completes, the Storage Lifecycle Policy will often have created the missing copy before the operator discovers the original failure.

6.2 Considerations when using Disk Storage

Disk storage that can only be presented to one Media Server (BasicDisk and AdvancedDisk storage using direct attached storage or SAN/NAS attached storage that cannot be zoned or presented to other servers) represents a single point of failure and, while it may be used as an initial storage target, the contents should always be duplicated to another medium as soon as possible.

6.3 Duplication of Backup Data outside of NetBackup Control

Duplication under NetBackup control involves creating a second copy of the data on a second storage destination and recording the existence of both copies within NetBackup. Each storage destination must be unique so duplication cannot, for example, occur between two tapes with the same label or the same disk on the same Media Server.

Duplication or replication of backup data can also occur outside of NetBackup control within the hardware storage layer. For example, a VTL may allow cloning of a virtual tape to either another

VTL or a physical media or a disk array supporting BasicDisk storage may allow its contents to a 'shadow' array at a different site. In both of these cases, although the data exists in two places, from a NetBackup perspective there is only one copy of the backup. This can present a problem if both copies are presented to the same NetBackup domain at the same time but can otherwise provide a useful aid to disaster recovery. The following sections describe how these features may be used for disaster recovery purposes.

Note: Although some OpenStorage appliances allow data to be replicated as part of the optimized duplication process, from a NetBackup perspective the replicated data is still treated as two distinct copies.

6.3.1 Considerations for replication within VTLs

Many VTLs support two types of replication that takes place outside NetBackup control:

1. Duplication of a virtual tape to a physical tape – in this case, provided the mapping between the virtual tape and physical tape is one-to-one (i.e. the physical tape has the same label and contents as the virtual tape), then the physical tape can be ejected from the tape library and sent to the secondary site.
2. Duplication of a virtual tape to another VTL in a secondary NetBackup domain – in this case, provided the secondary VTL is not visible to the primary NetBackup domain and the tape label ranges are unique across both domains, the replicated virtual tape can be used for disaster recovery at the secondary site provided that the virtual tape has been added to the EMM database in the secondary domain (e.g. by an inventory of the VTL on the secondary site).

In both cases, if the image records from the primary domain are available in the secondary domain (by means of partial catalog restore or replication), then it is a simple matter to restore backups using an alternative Media Server as described in section 3.2.2.4. If the image records from the primary domain are not available, then the tapes (virtual or physical) may be imported to create these records. VTL replication in this way can be used with both the traditional 'vault' disaster recovery model and the 'recovery without import' disaster recovery model (see sections 7.2 and 7.3) but should not be used in environments where a single NetBackup domain spans both the primary and secondary sites as NetBackup does not allow the same tape label to be used twice within one domain.

6.3.2 Considerations for replication within disk arrays

Disk arrays can be used to provide storage for BasicDisk storage units. In order to restore from a replicated BasicDisk volume at the secondary site, it is simply a matter of ensuring that the volume is mounted to the same mount point on the alternative Media Server configured as described in section 3.2.2.4 and that the image records are available.

If the primary and secondary sites form part of the same domain the storage at the primary site should be dismounted and replication suspended before the volume is mounted at the secondary site to avoid any possible access conflicts.

6.4 Encryption of data going off site

Best practice dictates that data that is sent off site on backup tapes should be encrypted. Various mechanisms for data encryption exist including NetBackup Client Encryption (in which data is encrypted on the client and must be restored to the same client), NetBackup Media Server Encryption Option (in which data is encrypted as it is written to tape), hardware encryption devices such as NeoScale and Decru, and tape drive encryption in LTO4 tape drives.

One important consideration when implementing encryption technologies is that the capability to decrypt the data must exist at the disaster recovery site for site disaster recovery to be possible. Some of the encryption technologies mentioned above (such as Media Server Encryption, and, with NetBackup 6.5.2 and onwards, tape drive encryption) use centralized key management in

White Paper: Implementing Highly Available Data Protection with NetBackup

which key information is located in a centralized database. These technologies offer simplified disaster recovery as the key management database can be copied between servers and backed up and recovered at a disaster recovery site.

Other technologies (such as hardware encryption and Client Encryption) use localized key management, making it difficult to recover data if the encrypting device or server is not available.

For clients using NetBackup Client Encryption, the encryption key path should be backed up using a non-encrypted backup and then restored to the disaster recovery client before the main (encrypted) backup is restored – this prevents the use of BMR for client recovery as the main backup is encrypted and cannot be restored until the key path is restored.

7.0 Disaster Recovery – Protecting against site failure

The term ‘disaster recovery’ usually means a major outage of some kind such as a site compromising incident rather than the loss of a single server, application or infrastructure component within a data center. This section looks at high availability at the site or data center level.

The configurations described here are typical examples and many variations on these configurations are possible offering varying degrees of protection.

7.1 Copy selection for restore

By default restore operations in NetBackup always use the primary copy (copy 1) for restore. In practice the primary copy of the backup is usually held locally for file, application and server recovery and the second copy of the backup data (copy 2) is sent to off site storage or the disaster recovery site. Before backups can be restored at the secondary site it is necessary to instruct NetBackup to use copy 2 for restores.

For versions of NetBackup prior to NetBackup 6.5 this can be done by using one of the following commands:

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

or:

```
bpimage -npc 2 -backupid <backup id>
```

The `bpchangeprimary` command is generally recommended for disaster recovery situations as it applies the change to all backups of a specific client in one pass and does not require knowledge of a specific backupid – it also supports the use of date ranges (see the NetBackup Commands Guides for more details). The `bpimage` command is more applicable to situations in which the primary copy of the backup is unreadable (damaged tape etc.) and only a specific backup needs to be changed. Both of these commands change information in the NetBackup image database and should be regarded as making permanent changes to the restore behavior.

From NetBackup 6.5 onwards there are two additional ways to select the copy to restore from:

1. Create a touch file `ALT_RESTORE_COPY_NUMBER` in the NetBackup root directory (`/usr/opensv/netbackup` or `<install path>\netbackup`) containing the copy number to be used for restores – this value is then applied to all restores for all clients until the file is removed.
2. Add the qualifier `-copy 2` to the CLI restore command `bprestore`

Neither of these approaches has any impact on the NetBackup image database and so can be regarded as only making a temporary change to the restore behavior.

7.2 Good – Dedicated Disaster Recovery Site with Vault

The simplest form of protection against site failure involves the storage of backups on tape media at an off-site location which may or may not be a data center that can be used for disaster recovery. Typically the ejected media are sent to a dedicated tape storage facility rather than a dedicated disaster recovery site and then transferred to the disaster recovery site when required.

The Vault Option in NetBackup can be used to manage the duplication, ejecting and tracking of tape media sent to an off-site location as well as ensuring that a current catalog backup accompanies the tapes.




Figure 2 - Disaster Recovery with Vault

Figure 2 shows a typical disaster recovery model using Vault. Backups are created at site A and are duplicated to a set of tapes that are then ejected from the library at site A and shipped to site B, a dedicated disaster recovery site where they are placed in a library. In the event of a site failure at Site A the catalog backup, made as part of the Vault processing at site A, is restored to the Master Server at Site B (which must have the same name as the Master Server at Site A) and, following some additional configuration steps to handle restore using different Media Servers, backups may then be restored to the disaster recovery servers at site B.

It may be argued that this approach provides the best disaster recovery solution as it ensures complete recovery of the entire NetBackup domain, however the cost of maintaining a dedicated disaster recovery facility that mirrors the production domain is prohibitive and, in most cases, some compromises are required in terms of the site configuration. These compromises lead to the need to reconfigure the domain when the catalog backup is recovered and adds to the complexity of the recovery process.

7.2.1 Benefits of a dedicated disaster recovery site

1. Using Vault to track the tapes sent to and from the off-site storage and the tapes required in the event of a disaster ensures that recovery.
2. As this solution uses a dedicated disaster recovery environment with its own Master Server, Media Server and tape library partition, regular disaster recovery test exercises can be carried out without impacting the production environment.
3. As full catalog recovery is carried out (including the BMR database), BMR can be used to recover servers easily at the disaster recovery site.
4. VTLs that support replication can be used to 'logically' off site virtual tapes.

7.2.2 Limitations of a dedicated disaster recovery site

1. This solution requires a dedicated disaster recovery environment with its own Master Server, Media Server and tape library partition. Although this may be a very small facility

located in another production data center it still represents equipment that cannot be used for other purposes.

2. The Master Server configuration for the disaster recovery environment must match that of the primary domain so if clustering is used for local resilience at the primary site a cluster must also be provided at the disaster recovery site (although this may be a single node cluster).
3. Tape media must be transported to and stored at off-site locations incurring transportation and storage costs.

7.3 Better (a) – Recovery without Import

The 'recovery without import' method of disaster recovery is a variant of the conventional disaster recovery approach which allows the primary and secondary sites to both operate as production domains. The name comes from the concept of 'importing' tapes into an existing NetBackup domain, essentially scanning the tapes and building the catalog entries from the information found on them in order to restore data from them. The 'recovery without import' model avoids the need to scan individual tapes by adding the image records to the database directly from a catalog backup.

The tapes, together with a catalog backup tape, are still sent to the secondary site but only the image database is restored from the catalog backup (the Vault Option can still be used to manage the duplication process and tracking of tape movements). The tapes associated with the primary domain are entered on the secondary domain as belonging to a media pool that is not used by any backup policy (and is not a scratch pool) in that domain to ensure they are not accidentally overwritten. (Depending on the library type and barcode schemes, bar code rules may be used to automatically add the tapes from the primary site into the appropriate media pool at the secondary site. Alternatively, if a known range of tapes is used, they could be manually added as part of the disaster recovery planning process.)

Once restored to the secondary site, the image records may be browsed and, because the tapes are known to the domain, restores may be initiated. The tapes themselves are not assigned within NetBackup at the secondary site, but because they are recorded in a 'private' media pool, they cannot be appended to or overwritten.

The following conditions must be met for 'recovery without import' to work successfully:

- The primary and secondary domain Master Servers must run the same operating system and version of NetBackup
- The barcode ranges used on tapes in the primary and secondary domains must be different
- The Media Server and client names in the primary and secondary domains must be unique
- Alternate restore Media Server settings, defined via the GUI or by setting `FORCE_RESTORE_MEDIA_SERVER` entries in the `bp.conf` of the secondary domain's Master Server, must exist to allow Media Servers in the secondary domain to restore backups created at the primary domain. These can be set during the initial configuration and left in place permanently, avoiding the need to set them up at the time of a disaster and allowing periodic disaster recovery practice exercises to be carried out.

In most cases a catalog backup is used to transfer the image information to the secondary domain, although it is possible to replicate the image data between the primary and secondary domains if UNIX or Linux Master Servers are used.

If an online catalog backup is used, the backup tape must be imported and then browsed to select the relevant client images rather than restored using the catalog recovery wizard.

Note: Although the catalog recovery wizard allows the recovery of only the 'NetBackup catalog image and configuration files', this option restores components such as the license database and, in the case of UNIX and Linux servers, the bp.conf file and should not be used to restore the catalog information to another Master Server unless suitable precautions have been taken to prevent these files from being overwritten.

If an offline catalog backup is used it must be configured to explicitly backup the individual paths below /usr/openv/netbackup/db, rather than backing up the whole path as one. In this case, only the path /usr/openv/netbackup/db/images should be restored to the Master Server at the secondary site.

The following diagram shows a configuration that uses replication of the image database rather than catalog backup and recovery.




Figure 3 - Recovery without Import using Image Database Replication

In this case the image database (/usr/openv/netbackup/db/images) on the Master Server at Site A is located on a separate mount point which is replicated to the Master Server at Site B. In the event of a site loss at the primary site, the replicated storage is mounted to the Master Server at Site B under a separate mount point from the main NetBackup image database and soft links are then used to map at the client level. (For example, the image database, /usr/openv/netbackup/db/images from the primary site is mounted under /primary/images on the secondary site Master Server and a link is created from /usr/openv/netbackup/db/images/ClientX to /primary/images/ClientX on this server to allow backups of ClientX to be browsed and restored).

7.3.1 Benefits of recovery without import

1. In this configuration the DR NetBackup domain can be another production domain as only the image records and media are being added to the environment rather than a full catalog recovery being carried out.

2. There is no requirement for the site configuration to be the same at both sites. Provided both Master Servers run the same operating system and version of NetBackup they can have different host names, as can the Media Servers within the two domains.
3. As the adding of tapes and recovering of image data into the secondary domain is relatively un-intrusive this approach allows frequent disaster recovery exercises to be carried out and, in some cases, may form part of daily operating routines – ensuring that the environment is always primed for rapid disaster recovery.
4. As only the image records are required at the secondary site it is possible to have a clustered primary Master Server and standalone secondary Master Server when cold catalog backup or replication is used.
5. VTLs that support replication can be used to ‘logically’ off site virtual tapes.

7.3.2 Limitations of recovery without import

1. Tapes must still be physically transferred between the primary and secondary sites using ground transportation. This means that it is possible that backup image records have been replicated to the Master Server at the secondary site, but the associated tapes have not yet been transferred to the site, so it is not possible to restore the most recent backups.
2. As the tapes are not assigned at the secondary site they do not expire and cannot be re-used. This may lead to media shortages following a prolonged outage. This can be addressed manually by determining which tapes have no valid images and moving them to the scratch pool, but this must be done with care to avoid accidentally overwriting valid backups.
3. Where replication of the image catalog is used the solution is limited to domains with a UNIX or Linux Master Server as it is not possible to present the NetBackup image database on a separate mount point in a Windows environment.
4. The indexing level used on the primary and secondary image databases must be the same – this should not be a problem as indexing is usually set to the default level of 2 in all environments.
5. As the full catalog backup is not restored in this case, BMR cannot be used for server recovery at the secondary site.

7.4 Better (b) – Media Server Promotion

The Media Server Promotion solution covers the specific use case where a single NetBackup domain spans two sites with a standalone Master Server at one (primary) site and at least one Media Server running on the same operating system at the other (secondary) site. A scheduled hot catalog backup is made (ideally to disk) on the Media Server at the secondary site on a regular basis (for example daily) and backup tapes are transferred between sites.

In the event of a failure of the primary site or the Master Server, the catalog backup is restored to the Media Server at the secondary site, which is then promoted to the role of Master Server.

This solution is suitable for small to medium enterprises that have two data centers with a minimum network connection of 100 Mbits/second between them. In a configuration of this type, tapes would still need to be moved between sites for off site storage, however where faster networks or wide area SAN connections exist, it may also be possible to duplicate the backups between sites and avoid the need to move physical media.

One key aspect of the promotion process used in this model is that rather than simply renaming the Media Server to be the Master Server, all of the existing attributes of the Media Server are preserved so there is no need to reconfigure the storage units or backup policies associated with that Media Server.




Figure 4 - Media Server Promotion

In the example above duplicate tapes created at Site A are transferred to site B and vice versa. In the event of a failure of Site B backups can be recovered at Site A. In the event of a failure at Site A the Media Server at Site B is promoted to the role of Master Server and backups from Site A are recovered at Site B.

7.4.1 Benefits of Media Server promotion

1. This solution offers a low cost disaster recovery capability without the need to use replication, clustering or a dedicated disaster recovery facility.
2. The solution also avoids the requirement for multiple NetBackup domains for production and disaster recovery.
3. As the full catalog backup is restored to the Media Server, BMR can be used to recover servers from the primary site to the secondary site.

7.4.2 Limitations of Media Server promotion

1. The Master Server and Media Server must be the same platform running the same version and patch level of NetBackup (NetBackup 6.0 MP5 or above).
2. Both servers must be standalone servers.
3. SPAS and NetBackup Access Controls (NBAC) cannot be used.
4. In the case of Windows servers, the install path (e.g. **C:\Program Files\VERITAS**) for the Master Server and Media Server should be the same.
5. Sufficient disk space must exist on the Media Server for a complete recovery of the catalogs to the same location that they resided in on the Master Server.

6. The Media Server must have the necessary components and services associated with the EMM database installed on it at the time that the switch takes place.
7. The Media Server must be powerful enough to assume the role of Master Server when needed.
8. The action of promoting the Media Server is relatively complex and, while the process is reversible, it does not lend itself to frequent disaster recovery test exercises.

7.5 Better (c) – Globally Clustered Replicated Master Server

One of the more time-consuming steps involved in a disaster recovery operation is the restoring of the NetBackup catalogs to the Master Server at the disaster recovery site and configuring the disaster recovery domain. Even with Vault management to ensure the correct catalog backup is available it is still necessary to recover the information and complete the configuration before restore of the mission critical servers can commence.

This phase of the recovery can be eliminated by replicating the catalogs to a standby node on the disaster recovery site. NetBackup allows this by establishing a global cluster between the two sites. In this configuration, the NetBackup catalogs are replicated between the sites under the control of the global cluster.




Figure 5 - Disaster Recovery Site with Replicated Cluster

Tapes are transferred to the disaster recovery site and loaded into the tape library in the same way that they are under Vault control in the previous examples, but the catalog recovery phase is no longer required as the Master Server is still the same, just running on a different physical node.

BasicDisk volumes may also be replicated between sites and mounted to a Media Server at the secondary site in the event of a disaster (see section 6.3.2).

In the event of a site loss at the production site, the Master Server cluster is simply failed over to Site B where a Media Server is available to handle restores. This Media Server is already

configured as part of the backup domain and the devices associated with it are known to NetBackup. The appropriate failover restore Media Server settings are also in place so that restore operations automatically use this Media Server when the Media Servers at Site A are unreachable.

As with the “Media Server Promotion” approach it is possible to continue making backups at the disaster recovery site and have these backups available when control is handed back to the primary site. The key advantage over the “Media Server Promotion” approach is that returning control to the primary site simply involved failing the cluster back when the primary site is available again.

7.5.1 Benefits of a replicated Master Server cluster

1. Global clustering simplifies the failover and failback processes ensuring minimal down time and risk of data loss during a site transition.
2. As the domain is simply deprecated and not completely lost during a site failover it is possible to continue making backups at the disaster recovery site with minimal configuration changes, making this model very suited to prolonged outages of the primary site.
3. Cluster failover and failback can be achieved relatively quickly making it easy to carry out frequent disaster recovery test exercises.
4. As all components of the NetBackup catalogs are replicated, BMR can be used to recover servers at the secondary site provided suitable BMR boot servers are available.

7.5.2 Limitations of a replicated Master Server cluster

1. The secondary site is effectively a disaster recovery site and the tape library at this site cannot be used for production backups. However, the use of library partitions and the creation of a second NetBackup domain at the secondary site would allow for the configuration to be mirrored, allowing each production site to support a disaster recovery component for the other site.
2. The separation distance between sites is determined by the replication technology in use and may be limited by certain replication technologies.
3. NetBackup security controls (SPAS and NBAC) are not supported for all combinations of NetBackup and Veritas Storage Foundation HA.

7.6 Best – dual site/single domain

Where sufficient site-to-site bandwidth exists the global clustering concept can be extended to provide a cross-site backup capability. The dual site/single domain model extends the concept of the globally clustered replicated Master Server to form a single domain with clients and Media Servers active at both sites under control of a common Master Server.

In this model not only are the NetBackup catalogs replicated to the alternate site, but the backups are also created on both sites (either via in-line copy or duplication depending on the configuration) so the loss of a single site does not represent a true ‘disaster’, simply the loss of a number of application servers. As the backup domain spans both sites, the loss of a single site simply results in a reduction of the backup and restore capability rather than destroying the backup environment.




Figure 6 - Dual site/single domain with stretched SAN

In figure 6 a wide area SAN connection allows Media Servers at both sites to write to tape drives and disk storage at both the local and remote sites. In this configuration backups may simply be written to tape or disk storage on the remote site or written to both sites using the in-line copy capability of NetBackup for improved resilience.

Duplication of backups can also take place off-line either over the SAN or, if no stretched SAN is available, over a dedicated backup network connection between the Media Servers at the two sites.

Configurations of this type rely on the presence of significant bandwidth between sites in either the network or SAN connections between sites. The introduction of space optimized, replicating storage appliances in NetBackup 6.5 allows a more efficient model to be adopted in which significantly less site-to-site bandwidth is required, as figure 7 shows.




Figure 7 - Dual Site/Single Domain with Optimized Duplication

7.6.1 Benefits of the Dual Site/Single Domain model

A configuration of this kind offers considerable advantages over more traditional configurations including:

1. There is no requirement for physical movement of media between sites because a copy of the backup is created on both sites.
2. Both sites may operate as full production data centers thus maximizing the usage of the infrastructure servers and storage.
3. Disk based backups may be used for recovery at the secondary site as well as tape backups.
4. As everything operates in a single domain, cross site backups are available instantly for restore at the alternate site, significantly reducing the RTO in the event of a site outage.
5. There is no true 'disaster recovery' step as both the NetBackup databases and the backup media exist at the secondary site. In the event of the loss of the primary site, recovery is simply a matter of restoring individual applications and servers at the secondary site.
6. As all components of the NetBackup catalogs are replicated, BMR can be used to recover servers at the alternate site provided suitable BMR boot servers are available.

7.6.2 Limitations of the Dual Site/Single Domain model

1. In most cases a single layer 3 network will not span two sites connected by a global cluster. This means that the virtual machine IP address will change when the cluster fails from the primary site to the secondary site, necessitating a DNS update before the Master Server is accessible again.

2. The separation distance between sites is determined by the replication technology in use and may be limited by certain replication technologies.
3. NetBackup security controls (SPAS and NBAC) are not supported for all combinations of NetBackup and Storage Foundation HA.

7.7 Disaster recovery options feature comparison

The following table compares the various approaches outlined in the previous sections against key capabilities and requirements for disaster recovery.

	Dedicated DR Site with Vault	Recovery without import	Media Server Promotion	Global Master Server Cluster	Dual Site/Single Domain
No requirement for dedicated DR servers	X	✓	✓	✓	✓
Supports BMR	✓	X	✓	✓	✓
Supports cross site duplication	X	X	✓	✓	✓
Supports BasicDisk replication	✓	✓	✓	✓	✓
Supports VTL replication	✓	✓	X	X	X
Supports Optimized Duplication	X	X	✓	✓	✓
Easy to test DR capability	✓	✓	X	✓	✓
Allows clustering of primary Master Server	✓	✓	X	✓	✓
Supports secure access (SPAS/NBAC)	✓	✓ ¹	X	✓ ²	✓ ²

¹ Uses SPAS/NBAC authentications of secondary domain

² Not supported for all combinations of NetBackup/Storage Foundation HA

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek
Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

Copyright © 2008 Symantec Corporation. All rights reserved. Symantec and the Symantec logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

1/08 13599373