



“ For over a decade Coolspirit have been supplying the UK’s top organisations with storage products and solutions so be assured we will meet your requirements head on.

It’s all about getting things right first time, quickly and simply! ”

Damon Robertson
Coolspirit Ltd

Our address

24 The Bridge Business Centre
Beresford Way
Chesterfield
S41 9FG

Get in touch

Call us on: 01246 454222
Email us: web@coolspirit.co.uk
Find us: [View location map](#)
Web: www.coolspirit.co.uk


Office hours

mon - thurs 8:30am - 5:30pm
fri 8:30am - 5pm
sat - sun Closed

“ Boost your storage buying power...
use ours! ”

Buy with confidence from
Coolspirit your authorised
Symantec Partner





Confidence in the connected world.

Backing Up VMware® with Veritas NetBackup™

George Winter | January 2009

Contents

- 1.0 EXECUTIVE OVERVIEW 3**
 - 1.1 INTENDED AUDIENCE..... 3
 - 1.2 GLOSSARY 3
 - 1.3 ADDITIONAL RESOURCES 4
- 2.0 BACKUP PARADIGMS – LOCAL AND OFF-HOST 4**
- 3.0 COMPARISON OF BACKUP METHODS 4**
 - 3.1 NETBACKUP CLIENT INSTALLED INSIDE VMWARE SERVICE CONSOLE..... 5
 - 3.2 NETBACKUP CLIENT INSTALLED INSIDE EACH VIRTUAL MACHINE 7
 - 3.3 PUREDISK CLIENT INSTALLED INSIDE EACH VIRTUAL MACHINE 8
 - 3.4 NETBACKUP FOR VMWARE – INTEGRATED WITH VMWARE CONSOLIDATED BACKUP (VCB)..... 10
- 4.0 IMPLEMENTATION, DEPLOYMENT AND BEST PRACTICES 11**
 - 4.1 CONFIGURATION 1 – NETBACKUP CLIENT INSTALLED INSIDE THE VMWARE SERVICE CONSOLE..... 11
 - 4.1.1 *Installation Procedure*..... 11
 - 4.1.2 *Configuration* 11
 - 4.1.3 *Configuring a NBU Policy*..... 12
 - 4.1.4 *Restoration procedure*..... 12
 - 4.1.5 *Hints, Tips and Best Practices* 12
 - 4.2 CONFIGURATION 2 – NETBACKUP CLIENT INSTALLED INSIDE EACH VIRTUAL MACHINE 12
 - 4.2.1 *Installation Procedure*..... 12
 - 4.2.2 *Configuration* 13
 - 4.2.3 *Configuring a NBU Policy*..... 13
 - 4.2.4 *Restoration procedure*..... 13
 - 4.3 CONFIGURATION 3 – PUREDISK CLIENT INSTALLED INSIDE EACH VIRTUAL MACHINE 14
 - 4.3.1 *Installation Procedure*..... 14
 - 4.3.2 *Configuration* 14
 - 4.3.3 *Configuring a PureDisk Policy*..... 14
 - 4.3.4 *Restoration procedure*..... 14
 - 4.3.5 *Hints, Tips and Best Practices* 14
 - 4.4 CONFIGURATION 4 – NETBACKUP FOR VMWARE, INTEGRATED WITH VCB 14
 - 4.4.1 *Installation Procedure*..... 14
 - 4.4.2 *Configuration* 15
 - 4.4.3 *Configuring a NBU Policy*..... 15
 - 4.4.4 *Restoration procedure*..... 16
 - 4.4.5 *Hints, Tips and Best Practices* 17
- 5.0 VMWARE BACKUP PROXY SIZING..... 18**
- 6.0 NETBACKUP 6.5.2 VCB INCREMENTAL BACKUP TECHNOLOGY EXPLAINED 20**

1.0 Executive overview

VMware® virtual infrastructure software is used by enterprises large and small to increase the efficiency and cost-effectiveness of their IT operations. Considered by Gartner to be a “mega-trend,” VMware is making its way into data centers of every size. Recognizing this trend, NetBackup™ has engineered innovative, award-winning data protection solutions designed specifically for VMware environments. This paper discusses best practices for designing solutions for and protecting VMware virtual machines.

As beneficial as virtual machine (VM) technology is, it also introduces new data protection questions and challenges. For example: Is it best to protect the virtual machine by backing it up via a NetBackup or PureDisk™ client? Is a NetBackup client inside the Service Console the answer? What about VMware Consolidated Backup? The advantages and disadvantages of these three backup configurations are discussed in detail in this document.

1.1 Intended audience

As there are many ways of using VMware technology, there are just as many methods available for protecting this innovative technology. System administrators and IT technologists can use this paper to determine one of three recommended solutions for protecting VMs. Each of these technologies has relative advantages and disadvantages.

1.2 Glossary

Backup proxy—System designated as the off-host backup system. At a minimum, the backup proxy needs to have the VMware Consolidated Backup framework software installed and at least a NetBackup client installed.

Converter Server—Originally designed for P to V and V to P conversions, the latest version provides integration with NetBackup for automatic VM ESX server registration from NetBackup backup images.

Guest OS—The operating system that runs on top of a VM.

Raw Device Mapping—An optional way to map physical SAN LUNs directly to a VM. Commonly used to enable application clustering and array-based snapshot technology.

RDM—See Raw Device Mapping.

Sync driver—Flushes OS buffers (Windows® only) before VMware Consolidated Backup snapshots are initiated. The sync driver is installed via VMware Tools. See also VSS Writer.

VCB—See VMware Consolidated Backup Framework.

Virtual machine—Software that creates a virtualized environment between the computer platform and its operating system so that the end user can install and operate software on an abstract machine. Note that the virtual machine designation does not imply any specific operating system version.

VM—An acronym for virtual machine.

VMDK—A designation specific to the files that comprise a VMware VM. These files are commonly called “vmdk” files because of the .vmdk extension that VMware adds to these files.

VMware Consolidated Backup Framework—An off-host backup API created by VMware. Designed to offload backup processing from the ESX server.

VMware Tools—Installed inside each VM. VMware Tools enhances VM performance and adds additional backup-related functionality.

VSS Writer—VMware replaced the sync driver with a Volume Shadow Copy Service writer beginning with the ESX 3.5 U2 release.

1.3 Additional resources

NetBackup for VMware configuration guide. For more information visit <http://support.veritas.com/docs/289771>.

Veritas NetBackup 6.5.2 Documentation Updates. Chapter 2 details how to configure and use the NetBackup for VMware feature set. For more information, visit <http://support.veritas.com/docs/302438>.

Veritas NetBackup 6.5.3 Documentation Updates. Includes detailed information specific to NetBackup for VMware. For more information, visit <http://support.veritas.com/docs/305408>.

Veritas NetBackup Snapshot Client Configuration. Discusses supported components in a NetBackup for VMware environment. For more information, visit <http://support.veritas.com/docs/288300>.

VMware Hardware Compatibility Guide. This is a web-based searchable guide that can provide compatibility information for systems, SAN, I/O devices, etc. For more information, visit <http://www.vmware.com/resources/compatibility/search.php>.

VMware SAN Configuration Guide. For more information, visit http://www.vmware.com/pdf/vi3_san_guide.pdf.

Understanding VMware Consolidated Backup. An introduction to the VMware Consolidated Backup technology. For more information, visit http://www.vmware.com/pdf/vi3_consolidated_backup.pdf.

Veritas NetBackup Backup Planning and Performance Tuning Guide. Provides significant detail related to NetBackup Media Server (and in-turn the backup proxy). For more information, visit <http://support.veritas.com/docs/307083>.

2.0 Backup paradigms—Local and off-host

Two backup paradigms for VMware are discussed in this document:

Local backup: These technologies involve installing a NetBackup or PureDisk client inside each VM or on the ESX Service Console. This backup methodology is popular because implementation is essentially the same as with physical machine backups. This process is described in configurations 1, 2, and 3.

Off-host backups: This design takes advantage of the VMware Consolidated Backup technology. Introduced with Virtual Infrastructure 3, this SAN-, NAS-, or iSCSI-based technology off-loads backup processing from the ESX server to a separate backup proxy server. NetBackup 6.5.1 adds Granular File Restore from full VM (vmdk) backups—an award-winning technology only offered by NetBackup. This technology is described in configuration 4.

3.0 Comparison of backup methods

There are a number of ways of protecting VMs. In this paper, we cover four of the most popular ways of protecting VMware. Table 1 provides a high level overview of each of these technologies that can be useful for determining which technology best suits the needs of your specific VMware environment. Table 2 provides guidelines related to performance and hardware requirements.

Recommended For	NBU Client in Service Console	NBU Client in Virtual Machine	PureDisk Client in Virtual Machine	NBU for VMware VCB Integration
Simplified disaster recovery (vmdk)	✓	X	X	✓
Individual file restore	X	✓	✓	✓ (Windows)

Recommended For	NBU Client in Service Console	NBU Client in Virtual Machine	PureDisk Client in Virtual Machine	NBU for VMware VCB Integration
Backup over LAN	✓	✓	✓	✓ (NetBackup 6.5.2 +)
Backup over SAN / iSCSI	X	X	X	✓
Individual file search (Windows)	X	✓	✓	✓
Incremental backups	X	✓	✓	✓ (NetBackup 6.5.2 +)

Table 1: Solution comparison

	NBU Client in Service Console	NBU Client in Virtual Machine	PureDisk Client in Virtual Machine	NBU for VMware VCB Integration
ESX server resource efficiency	● ○ ○ ○ ○	● ○ ○ ○ ○	● ● ● ○ ○	● ● ● ● ○
Backup performance (speed)	● ○ ○ ○ ○	● ○ ○ ○ ○	● ● ● ○ ○	● ● ● ● ●
Additional hardware requirements	● ○ ○ ○ ○	● ○ ○ ○ ○	● ● ○ ○ ○	● ● ● ○ ○
● ○ ○ ○ ○ = Least ● ● ● ● ● = Most				

Table 2: Performance comparison

The following sections examine each of the four backup methods described in Table 1 and discuss the advantages and disadvantages of each method.

3.1 NetBackup client installed inside VMware Service Console

Reasonably simple to implement, this could be considered an off-host backup technology in the sense that no NetBackup software is installed inside the VM (Figure 1). However, unlike off-host backup technologies, in this configuration backup impact occurs on the ESX server. Installing the NetBackup client inside the Service Console gives direct access to the files that make up the VMs—the vmdk files. It should be noted that VMware has indicated that at some point the ESX Service Console will no longer be available. This should be a consideration for determining if this backup method should be selected.




Figure 1: Client in the ESX Service Console

This method is easiest to use if the VMs are powered off. In this state, the VMs are static and unchanging. If the VMs are powered on, additional pre-backup processing or scripting is recommended to ensure that the VMs are in a consistent state during backup operations. Implementing this would involve using the ESX server's built-in snapshot functionality. For more information, see section 1.3.

The relative advantages and disadvantages of this backup configuration are described as follows:

Advantages:

- Installation is clear-cut. A NetBackup client can be easily installed on the VMware Service Console. Configuring a NetBackup policy using this technology is as straightforward as any standard client backup policy configuration.
- Entire VM restores are simple. Client inside the Service Console provides backup and restore access to the VM (vmdk) files.
- The Service Console OS can also be backed up using this method.
- Back up the entire VM by backing up the vmdk files.
- Supports LAN and SAN implementations.

Disadvantages:

- Backup processing on one VM impacts the system resources available to all remaining VMs located on the ESX server.
- Script creation and maintenance is likely to be required if consistent backups of live VMs is required.
- Single (OS level) file restores are not possible directly from the NetBackup interface. However, single file restores can be performed using VMware created tools such as "mountvm.exe". For more information, see the resources listed in section 1.3.

- Does not support database or application backups.
- Not supported with ESX 3i (ESXi) implementations.

3.2 NetBackup client installed inside each VM

In spite of the virtualization technologies involved, VMs are complete OS installations hosted on virtualized hardware. These installations can be backed up the using the same basic techniques as their physical counterparts—with a NetBackup client inside the guest OS (Figure 2). Running a NetBackup client inside the VM is supported. Standard OS support rules apply. Backing up a VM in this way is essentially the same as backing up a physical machine.




Figure 2: NetBackup Client in the VM

Advantages:

- Simple and familiar implementation. Traditionally, most physical machine backups have been performed this way, making the transition to VM backups using this technology a straight-forward task.
- Single file backups are supported.
- All backed up data is correctly referenced in the NetBackup catalog to the originating VM.
- Restoration directly to the VM is supported.
- Incremental backups are easily configured.
- Advanced backup technologies such as synthetic backups are supported.

- Database backups are supported as well. Configuration is as simple as installing and configuring the appropriate database agent.

Disadvantages:

- Resource intensive backups often place a heavy load on shared network and CPU resources located on each ESX server.
- Entire VM (at OS level) restores are more complex and time-consuming.
- Simplified disaster recovery (restoring at the vmdk level) is not available.
- The backup processing load on one VM will negatively impact system resources available to other VMs hosted on the same physical server. Backup scheduling should take this issue into account.
- Client software installed inside each VM needs to be maintained and updated.
- The VM must be powered on for backup processing to occur.

3.3 PureDisk client installed inside each VM

In virtual environments, a traditional streaming backup with a standard client in the VM puts a tremendous strain on the environment because it consumes significant CPU and network bandwidth in the virtual infrastructure.

PureDisk addresses these typical challenges by offering a disk-based backup solution that dramatically reduces the size of backups and the network bandwidth required to perform them by using a data deduplication technology at the source (or client) of a backup, which in this case would be within the VM itself. PureDisk identifies redundant sub-file data segments and only sends and stores unique segments on a global basis. As a result, PureDisk eliminates the need to send or store duplicate files, such as OS files and executables, as well as redundant segments of files or databases as they are edited or changed over time. PureDisk also offers centralized web based management of all backups and supports most major operating systems in addition to supporting databases like MS® Exchange and SQL Server®.

This architecture essentially provides a solution that closely resembles the backup of physical servers, is simple to configure, and provides the same file level restore capabilities of a traditional backup.




Figure 3: PureDisk Client in each VM

Advantages:

- Simple and familiar implementation. Traditionally, most physical machine backups have been performed this way, making the transition to VM backups using this technology a straight-forward task.
- Single file backups and restores for all OSs are supported.
- All backed up data is correctly referenced in the PureDisk metabase (catalog) to the originating VM.
- Restoration directly to the VM is supported.
- SQL Server and Exchange backups are supported as well. Configuration is as simple as configuring the appropriate type of backup.
- Extremely light, efficient data transfer and storage. Storage for backups can often be reduced by 10x–50x. Bandwidth consumption for daily full backups is reduced by up to 500 times.
- Lower overall resource utilization on the ESX console.

Disadvantages:

- With this architecture the full VM image (vmdk files) is not backed up making full VM (at OS level) restores more complex.
- Disaster recovery restores are complicated and time consuming. Simple restores based on vmdk files are not possible.
- The backup processing load on one VM may negatively impact system resources available to other VMs hosted on the same physical server. We recommend that backup scheduling be

planned to take this issue into account. For example, the number of simultaneous VM backups should be limited.

- Client software installed inside each VM needs to be maintained and updated.
- The VM must be powered on for backup processing to occur.

3.4 NetBackup for VMware—Integrated with VMware Consolidated Backup

NetBackup for VMware provides an alternate-client backup technology for VMs with integrated file recovery capabilities (Figure 3). In conjunction with VMware Consolidated Backup off-host data protection technology, NetBackup 6.5 builds on and significantly enhances VMware Consolidated Backup. Basic VMware Consolidated Backup integration is included with NetBackup 6.5. Enhanced Granular File Restore from a vmdk backup was originally released with NetBackup 6.5.1 and has been enhanced for NetBackup 6.5.2 (and later) with additional features such as incremental backups. This paper focuses on and describes features that are available in the NetBackup 6.5.2 (and later) releases.

VMware Consolidated Backup is a VMware created backup API that provides a centralized backup facility. VMware Consolidated Backup is an off-host backup technology that leverages a centralized proxy server (for example, backup proxy), which in turn significantly reduces the backup processing load on production VMware ESX hosts.




Figure 4: NetBackup with VMware Consolidated Backup

Advantages:

- Two restore options are available from a single backup; single file restore or entire (vmdk level) VM restore.
- NetBackup integration with the VMware Converter technology greatly simplifies restores of VMs (extremely useful in disaster recovery scenarios).

- All backed up data is correctly referenced in the NetBackup catalog to the originating VM.
- Backup impact on the target VM and other VMs hosted on the same physical machine is minimized.
- No NetBackup client software needs to be installed on the VM or inside the ESX Service Console. Optionally, client software can be installed to assist with file level restores.
- Integrated with VMware Virtual Center to provide easy VM discovery.
- Distributed Resource Scheduling (DRS) and VMotion aware.
- File level incremental backups are now possible (requires NetBackup 6.5.2 or higher and VMware Consolidated Backup Framework 1.1 or higher—Windows only). This can significantly reduce backup times, data stored on disk/tape, and backup impact.
- Supports SAN, iSCSI and NAS environments. Datastores created on NAS require NetBackup 6.5.2 or later and the VMware Consolidated Backup Framework 1.1 or later.
- With the NetBackup 6.5.3 release, a VMware provided VSS Writer is supported within the VM. VSS Writer can provide basic application (and database) backup functionality by quiescing the application before the VMware Consolidated Backup snapshot is created. VSS support is only available when using the ESX 3.5 U2 or later release as well as the VMware Consolidated Backup Framework 1.5 or later.

Disadvantages:

- A VMware backup proxy is required. This can be hosted on a Windows 2003 or 2008 server and in many cases can be hosted on the same server as a NetBackup media server.
- Live (hot) application or database backups require additional pre- and post-backup processing to ensure backup data consistency.
- Granular file restore and incremental backups are currently available only for Windows VMs.

4.0 Implementation, deployment and best practices

The following sections describe the configuration, installation and set up of the backup components for each of the four backup methods.

4.1 Configuration 1: NetBackup client installed inside the VMware Service Console

4.1.1 Installation procedure

There are two major components of installing this configuration:

The first installation task is the easiest as it simply involves the installation of a NetBackup client inside the Red Hat Linux® based ESX Service Console.

The second task is technically not required but is recommended. This involves implementing a script to prepare the VM for backups. If the VM is powered on and running, the data in the VM files (vmdk) files will be in an inconsistent state and changing. The VM (vmdk) files can be backed up; however, restores are not assured as the backup is considered crash-consistent. To avoid this issue, scripting would be required to either shut down the VMs or, more commonly, create a snapshot of the VMs. NetBackup does not provide scripts for this purpose.

4.1.2 Configuration

There is no specific hardware configuration required for this deployment. There are a few simple requirements that are similar to any standard network backup configuration. For example, the ESX

Service Console must have network access to the NetBackup media server and its hostname must be resolvable.

4.1.3 Configuring a NetBackup policy

While NetBackup policy configuration is reasonably straightforward, a few issues should be kept in mind.

When defining backup selections, the base VMFS directory on the ESX server should not be used as a backup selection. Instead, a child folder within the vmfs directory should be selected. For example, to back up VMs that exist within the “datastore1” datastore, you would select the following path on the ESX server:

```
/vmfs/volumes/datstore1
```

as the backup path within the Backup Selections tab on the NetBackup policy definition.

4.1.4 Restoration procedure

Using this procedure, the VM vmdk files have been safely backed up. Restoring the VM based off of these vmdk files is a two-step process. The basic process is as follows:

1. Using the NetBackup Backup Archive and Restore GUI, the vmdk files should be restored to either space located on the ESX Service Console or an NFS mount that is accessible to the ESX server.
2. Once the vmdk files have been restored, they must be reintroduced or registered to the ESX server. This is done via the `vcbRestore` command. The entire procedure is outlined in the *Virtual Machine Backup Guide*. For more information, see section 1.3.

4.1.5 Hints, tips and best practices

There are a number of areas where configuring hardware and software can greatly enhance the reliability of VM backups as well as decrease the impact backups have on VMs.

- Limit the number of simultaneous backups that occur on a datastore. This limits the impact that backup operations have on each datastore and in turn, decreases the impact that backups will have on all VMs that share that datastore. This can be configured in the policy definition using the Limit Jobs Per Policy attribute (see Figure 4).

Configure (align) NetBackup policies with VMware datastores. To do this, create each policy so that every VM defined in a policy resides on the same datastore.

- Configure the policy to limit the number of jobs that run simultaneously, therefore limiting the number of backup operations that occur against this datastore.

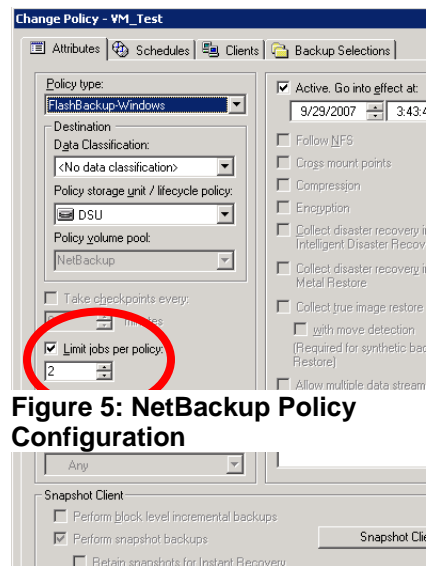


Figure 5: NetBackup Policy Configuration

4.2 Configuration 2: NetBackup client installed inside each VM

4.2.1 Installation procedure

From an installation perspective, this configuration is one of the most straightforward. A NetBackup client is simply installed inside the VM. The installation procedure for a VM is essentially the same as if the OS were hosted on a physical (not virtual) machine.

4.2.2 Configuration

For standard file backups, client configuration in the VM is the same configuration procedure as for a physical machine. Simply follow the installation instructions provided with the client.

4.2.3 Configuring a NetBackup policy

Basic NetBackup client policy configuration should be modified to take into account the physical layout of the ESX server, the VMs targeted for backup, and the datastore or datastores that the VMs reside on. The VM backup policies should be scheduled so that a minimum number of VMs backups occur simultaneously.

4.2.4 Restoration procedure

One advantage of this configuration is that restores can easily be and are typically directed toward the original VM. The previously installed NetBackup client makes this process exactly the same as if the restore were occurring to a physical host.

Hints, tips and best practices

- This backup configuration is an on-host style of backup. As such, backup activities on a single VM can create a significant amount of load on the parent ESX server and therefore indirectly impact every other VM that is also hosted on the ESX server. Backup policies should be defined to limit the number of simultaneous backup jobs that are running on each physical ESX server.
- NetBackup Synthetic Backup technology can be used to minimize or eliminate full backups. By only performing incremental backups, only the amount of data that has changed since the previous backup would be copied to the NetBackup media server, significantly decrease the I/O associated with backups and the amount of backup network traffic.
- On Windows hosts, the Windows Change Journal can also be implemented to reduce the backup impact that occurs during incremental backups.
- This backup technique lends itself very well to database backups. Configuring a database backup in a VM is essentially the same as configuring the same database backup on a physical machine. This technique can simplify and enhance VM database backups, often times providing incremental backup capabilities and restores directly to the VM.
- Limit the number of simultaneous backups that occur on a datastore. This limits the impact that backup operations have on each datastore and, in turn, decreases the impact that backups will have on all VMs that share that datastore. This can be configured in the policy definition using the Limit Jobs Per Policy attribute (see Figure 5).
- Configure (align) NetBackup policies with VMware datastores. To do this, create each policy so that every VM defined in a policy resides on the same datastore.




Figure 6: NetBackup Policy Configuration

4.3 Configuration 3: PureDisk client installed inside each VM

4.3.1 Installation procedure

From an installation perspective, this configuration is one of the most straightforward. A PureDisk client is simply installed inside the VM. The installation procedure for a VM is essentially the same as if the OS were hosted on a physical (not virtual) machine.

4.3.2 Configuration

Client configuration in the VM is the same configuration procedure as for a physical machine. Simply follow the installation instructions provided with the client.

4.3.3 Configuring a PureDisk policy

Basic PureDisk client policy and dataselection configuration should be modified to take into account the physical layout of the ESX server, the VMs targeted for backup, and the datastore or datastores that the VMs reside on.

4.3.4 Restoration procedure

One advantage of this configuration is that restores can easily be directed toward the original VM, and typically are. The previously installed PureDisk client makes this process exactly the same as if the restore were occurring to a physical host.

4.3.5 Hints, tips and best practices

- This backup configuration is a local (on-host) style of backup. Backup policies should be defined to limit the number of simultaneous backup jobs that are running on each physical ESX server.
- The PureDisk deduplication technology will result in only new unique data being backed up. This will significantly decrease the I/O associated with backups and the amount of backup network traffic generated. This will limit the impact backups have on all VMs hosted on this ESX server.
- Limit the number of simultaneous backups that occur on a datastore. This limits the impact that backup operations have on each datastore and in turn decreases the impact that backups will have on all VMs that share that datastore.
- Configure (align) PureDisk dataselections and policies with VMware datastores. To do this, create each dataselection and policy so that every VM defined in a dataselection or policy resides on the same datastore.
- Limit the amount of dataselections and policies as much as possible.

4.4 Configuration 4: NetBackup for VMware, integrated with VMware Consolidated Backup

4.4.1 Installation procedure

Configuration and installation of NetBackup for VMware is relatively straightforward. The following is a summary of the installation steps required to implement VMware Consolidated Backup with NetBackup 6.5.2 or later (Figure 6). For more information, see the *NetBackup 6.5.2 Documentation Updates* guide.

Step 1. Ensure that the hardware (especially the SAN or iSCSI environment) is configured properly. The datastore where the target vmdk files exist must be “visible” and accessible to both the ESX server and the backup proxy. VMware also has specific hardware and configuration requirements. An important component of reliable VMware Consolidated Backup backups is a properly configured SAN/iSCSI environment. For more information about VMware SAN requirements, see the *VMware SAN Configuration Guide* referenced in section 1.3.

Note that with the 6.5.2 and later release, NAS based VMware datastores are now supported (requires VMware Consolidated Backup Framework 1.1 or later).




Figure 7: VMware Consolidated Backup

Step 2. On the backup proxy, install the supported version of the VMware Consolidated Backup Framework. The installation of this VMware component is straightforward and requires no specific configuration tasks.

Step 3. Install the NetBackup component of choice (master server, media server or enterprise client) on the VMware backup proxy. As a best practice, we recommend that the backup proxy be configured as a NetBackup media server. For more information, see the “Hints, tips and best practices” section later in this document.

As you can see, once the hardware is installed and properly configured, the software installation procedure is straightforward.

4.4.2 Configuration

There are two main configuration tasks that need to be performed with this backup method:

1. Configure the VMware Consolidated Backup components within NetBackup. This refers to the Virtual Center Server (or, optionally, the individual ESX servers if no Virtual Center Server is used).
2. VMware backup proxy. This is the NetBackup component where the VMware images are mounted and backed up. This system can be configured as a NetBackup master server, media server or enterprise client. However we recommend that this system be a NetBackup Media Server. Besides the NetBackup software component, the VMware Consolidated Backup Framework is installed on this system as well.

Both of these entities are defined in more detail within the NetBackup Admin Console. For more information, see the *NetBackup for VMware configuration* guide.

4.4.3 Configuring a NetBackup policy

A NetBackup for VMware policy can either be manually created or created using the NetBackup Snapshot wizard. With NetBackup 6.5.1 (or later), integration with either the ESX server or the Virtual Center Server provides simplified VM discovery.

There are several NetBackup for VMware policy attributes that are specific to VMware backups. Optimal backup performance can be achieved if these attributes are correctly applied.

A description of these attributes and recommended settings are as follows:

Client definitions: This policy attribute is probably the single most important policy setting in VMware environments. When adding clients to a policy, we recommend that clients be aligned with the storage or

datastore on which these VMs reside. For example, if you have 40 VMs configured on two datastores and 20 VMs reside on each datastore, two policies could be created. Each policy would be aligned with the 20 VMs that are installed on their respective datastore. This technique would allow for direct control over how many simultaneous backups and subsequent VMware Consolidated Backup snapshots that could be created at any given time. This would also allow the backup administrator to limit that impact that snapshot and redo creation and deletion would have on each individual datastore, improving backup reliability and minimizing I/O contention. This policy definition recommendation can be extended for every VM client/datastore relationship in the environment.

Snapshot Mount Point: This is defined within the Snapshot Method Options portion of the policy definition (Figure 7). Only one snapshot mount point can be defined per policy, but multiple policy definitions can be used to take advantage of multiple snapshot mount points. While the size of this mount point must be at least as large as the total number of simultaneous VM backups, the I/O performance of this mount point is an important consideration as well. When the FullVM backup technology is used, the VM vmdk files are copied to this mount point. The more I/O capacity this mount point has, the faster the vmdk files can be copied to this mount point and the quicker the VM snapshot

can be released. This limits the amount of time that the snapshot exists which in turn minimizes the amount of redo log creation. All of this mitigates the I/O impact that backup operations have on the datastore.

Limit Jobs Per Policy: This policy attribute can be used to limit the number of simultaneous snapshots that occur on each datastore. Different storage environments can support a varying number of simultaneous backups. The recommended number of simultaneous jobs that are configured for each policy/datastore combination will depend on the quality of the storage infrastructure involved. Best practices would dictate that this number should be in the low single digits. One technique for determining the maximum number of simultaneous supportable snapshots is to run test backups, gradually increasing the Limit Jobs Per Policy attribute until policy creation and deletion no longer reliably occurs, or until the I/O impact on the VMs associated with the datastore negatively impacts the performance of the associated VMs.

Client Schedule and Incremental Backups: NetBackup 6.5.2 introduces file level incremental backups based on vmdk-based granular file restore (Windows only) full backups. This is made possible because NetBackup is uniquely able to index, search for, and restore individual files from a vmdk backup. Using this vmdk backup as the basis for a full backup, NetBackup can compare the files that exist at the time of this full backup and perform subsequent, low impact file-level incremental backups. This process can significantly reduce backup times, I/O impact on the ESX datastore and disk/tape storage—all without limiting restore options. Regardless of whether a full or incremental backup has been performed, any version of any file that existed at backup time can be restored. For more information related to this configuration, see section 6.0.

4.4.4 Restoration procedure

Single file restores: NetBackup for VMware does not require that a NetBackup client be installed inside either the VM or the VMware Service Console. For this reason, direct restores to the VM are not possible unless a NetBackup client is installed inside the VM. Alternatively, an alternate client restore can be performed to a Windows share, and the restored files accessed and transferred to the VM through this share.




Figure 8: Backup Proxy Snapshot Mount Point

FullVM restores: Entire VM restores can be manually performed by restoring the vmdk files to a staging area and then moving and registering the vmdk files to the target ESX server using the VMware vcbRestore.exe command. Alternately, automated full VM restores can be performed by directing FullVM restores to a VMware Converter server. This entire restore procedure is completely automated by NetBackup. As part of the restore process, VMware Converter automatically restores the VM to the specified ESX server and automatically registers the VM. Once this fully automated process is complete, the VM is completely ready for use.

4.4.5 Hints, tips and best practices

- The success of VCB snapshot creation and deletion can be directly attributed to two things: (1) the amount of I/O occurring on the VM datastore during snapshot creation, and (2) the correct design of the I/O substructure associated with each datastore. To avoid snapshot-associated issues, backups should be scheduled during times of relative low I/O activity on the VM. Reducing the number of simultaneous backups (and, in turn, simultaneous VMware Consolidated Backup snapshots) can help with this as well. This can be done via the Limit Jobs Per Policy attribute within the policy definition. For correct I/O design and implementation, consult the VMware created documentation listed in section 1.3.
- Create and align the NetBackup policy with each VMware datastore (storage). For example, 10 NetBackup for VMware policies would be created in an environment that has 10 datastores. This provides the ability to manually isolate and control the amount of backup-related I/O that occurs per datastore. In this way the amount of backup-related I/O can be controlled and, in turn, limit the impact that backup I/O has on the target VMs.
- The VMware backup proxy can be configured as a NetBackup master server, media server or enterprise client. We recommend that the backup proxy be configured as a media server. The backup proxy is a natural focal point of backup-related I/O. Backup proxy access to VM files is typically made through a fast Fibre or iSCSI connection. With the NetBackup 6.5.2 release, a NAS-based datastore is supported as well. It makes sense in most configurations to avoid configuring the backup proxy as a NetBackup enterprise client as this forces all of the VM backup data through the NetBackup client network, typically a slow, shared resource.
- The use of multiple backup proxy servers is supported with NetBackup. Once a single backup proxy is saturated with backup processing (for more information, see section 5.0), another backup proxy can be added to increase backup throughput capacity.
- Once a VMware Consolidated Backup snapshot is created, data is transferred from the VM datastore to the backup proxy mount point. The completion speed of the snapshot process can be significantly enhanced if care is made to ensure that the data path from the datastore to the snapshot mount point (that is, staging area) is as fast as possible. The snapshot mount point should be configured over as many dedicated spindles as possible. While we recommend that the snapshot mount point be configured as fast disk, it does not need to be highly available. The snapshot mount point is only used temporarily during the backup process and is not a permanent repository of backed up data. The vcbMounter.exe command can be used to perform snapshot creation and transfer performance tests. This can be done by invoking the vcbMounter.exe command from the backup proxy. For example, if you wanted to test the snapshot throughput rate of a VM named vm100.veritas.com to a backup proxy named proxy1.veritas.com, you would run the vcbMounter.exe command from the backup proxy as follows:

```
vcbMounter.exe -h esx1.veritas.com -u root -p foobar -a  
ipaddr:vm100.veritas.com -r d:\mnt\vm100.veritas.com -t fullvm
```

Where the following VMware Consolidated Backup components are defined as:

ESX Server = **esx1.veritas.com**

Root user on esx1.veritas.com = **root**

Root password = **foobar**

Virtual Machine = **vm100.veritas.com**

Mount point on Backup Proxy = **d:\mnt\vm100.veritas.com**

- Multiple mount points (used to stage VMware Consolidated Backup snapshots) can be defined and used on the backup proxy server (Figure 8). Properly configured, additional VMware Consolidated Backup mount points (holding tank) can increase data backup throughput on the backup proxy by using additional staging areas to extend I/O capacity. New staging mount points that are defined should be created on separate and dedicated I/O channels, isolating I/O traffic from other mount points that exist on a given backup proxy.
- Do not make the staging mount point larger than necessary. If possible, keep the size of the staging mount point less than 1 TB.
- Each mount point (used to stage VMware Consolidated Backup snapshots) should be a separate file system created on a dedicated SCSI or Fibre channel bus. However, if the mount point shares space with other files or data, the mount point file system should be defragmented at regular intervals, ensuring maximum file system performance during the transfer of data from the ESX datastore to the backup proxy. This is especially important when FlashBackup™ style backups are performed.
- Separate HBAs (host bus adapter) should be used for each I/O path on the backup proxy. For example, for the destination storage unit (disk/tape), each staging mount point and the connection to the datastore(s) should all be configured on separate HBAs. This ensures that there is no I/O contention on the backup proxy. We recommend that each HBA be located (if possible) on separate internal buses within the backup proxy server.
- Upgrade to the latest version of VMware Virtual Infrastructure. This includes the latest version of ESX server, Virtual Center Server and VMware Consolidated Backup Framework. Newer versions of Virtual Center components typically have enhancements that improve VMware Consolidated Backup snapshot performance and reliability.
- When performing incremental backups (NetBackup 6.5.2/6.5.3), both cumulative and differential incremental backups are supported. Restoring an entire VM to a specific point in time involves first restoring the vmdk file, booting the VM, and then restoring subsequent incremental backups. The number of restores required to perform this restore can be limited to two. This is accomplished by employing the cumulative incremental backup type. Cumulative incrementals back up all of the data that has changed since the previous full backup. In this case, the previous full backup is a vmdk Granular File Restore backup.

5.0 VMware backup proxy sizing

During VMware Consolidated Backup backups, the backup proxy server performs a significant amount of backup processing. Proper sizing of the backup proxy server can help ensure maximum backup performance of the VM environment. Accurately characterizing the capacity of the backup proxy can be broken down into three main areas:

1. VMware Consolidated Backup data path

This is the entire data path that the VMware Consolidated Backup created data will follow during the backup lifecycle. During the first phase of the backup process, the vmdk files are transferred from the ESX server datastore to the VMware Consolidated Backup staging area for a FullVM style backup. This occurs as follows (see Figure 7):

- ① After the snapshot creation, VM data flow occurs between the ESX datastore and the backup proxy. Both iSCSI and Fibre channel technologies are supported for this connection.
- ② Data is then transferred from the backup proxy to a disk staging area where backup processing occurs.




Figure 9: VMware Consolidated Backup Data Path

Remember that both data paths (1 and 2 in Figure 7) described above are used during the initial copy of the vmdk file. The throughput of the *entire* data path described here is an important component when determining backup performance. The actual vmdk backup does not begin until the entire file copy process (from the datastore to the staging area) is completed. If network (NAS) based backups are configured, the performance of this initial copy process will be limited to network throughput speeds.

2. NetBackup data path

With VMware Consolidated Backup backups, the entire backup process is not over until the process of writing data to the NetBackup storage unit is complete. This is why the path from the VMware Consolidated Backup staging area to the NetBackup storage unit should be a design concern.




Figure 10: NetBackup Data Path

Once the vmdk files are copied to the backup proxy the next step is for them to be backed up by NetBackup. The backup proxy can be a NetBackup master server, media server or enterprise client. To maximize backup throughput, the backup proxy should be configured as a master or media server so that client data is written directly to a DSU or tape (Figure 8) and not first sent over the network, typically a relatively slower transfer medium than storage directly attached to the NetBackup media server.

3. Backup proxy sizing

The VMware Consolidated Backup process offloads backup processing from the ESX server to the backup proxy. This means that special care should be taken to ensure that the backup proxy has enough system resources to support the amount of I/O that will be required of it. Here are some basic guidelines that should be followed when designing the backup proxy:

- Backup proxy server I/O performance is generally more important than CPU performance.

- CPU, I/O, and memory expandability should also be a consideration when choosing a server.
- Size the CPU to support 10 MHz of CPU available per 1 MB/sec of data throughput in **and** out of the backup proxy.
- The internal bus of the backup proxy should be fast enough to support the I/O devices connected to it. If multiple I/O ports are used, a system with multiple internal buses should be considered to support the additional I/O.
- For more information about NetBackup sizing, refer to the *Veritas NetBackup Backup Planning and Performance Tuning Guide*.
- The VMware Consolidated Backup snapshot mount point (staging area) should be sized using the following equation:

$$\text{Mount Point Size (GB)} = (\text{NUM_VM}) \cdot (\text{AVG_SIZE})$$

Where:

NUM_VM = Largest number of VMs to be backed up simultaneously

AVG_SIZE = Average size of the largest VMs to be backed up simultaneously.

For example, if 5 VMs are to be backed up simultaneously, take the 5 largest VMs in the environment, calculate their average size and use that number (as **AVG_SIZE**) in this equation.

- Remember that the VMware Consolidated Backup snapshot area is only a temporary staging area and is automatically maintained and reused as VM backups are processed. While an extremely small staging area can create a backup bottleneck in the VMware Consolidated Backup process, if the backup proxy staging area is designed large enough to support a nominal amount of VMs during the backup process, the staging area size should not impede backup performance.
- I/O throughput is another performance metric that should be considered when sizing the backup proxy. For example, a 2 Gb Fibre connection should be able to transfer either a VMware Consolidated Backup snapshot or backup data at a nominal transfer rate of 140 MB/sec. For more information refer to the *Veritas NetBackup Backup Planning and Performance Tuning Guide*.
- A backup proxy sizing calculator can assist with the proper selection of hardware for the backup proxy. This calculator is available through your Symantec systems engineer.

Conclusion:

Overall backup performance of each backup proxy will be defined by the slowest component of the entire backup data path. These components are:

- Backup proxy system resources: CPU, internal bus, RAM
- VCB snapshot creation time
- Size of the staging area located on the backup proxy
- I/O performance of data path between ESX datastore and backup proxy staging area
- Data path between backup proxy staging area and the NetBackup storage unit

6.0 NetBackup 6.5.2 VMware Consolidated Backup incremental backup technology explained

NetBackup 6.5.2 provides vmdk file backups that are indexed and cataloged (Windows only). This process makes possible the ability to restore individual files (Word® documents, Excel® spread sheets,

etc.) without needing to first restore the entire vmdk file. During this process NetBackup also stores detailed information for every file, including creation and modification time. This allows NetBackup to perform subsequent incremental backups by comparing the contents of a VM (the full backup) against the stored catalog of the initial full (vmdk) image. If changes are detected, the VMware Consolidated Backup file level image is used to back up only the files that have change since the vmdk full backup. This technology is unique to NetBackup it is the only backup product that can index vmdk files at backup time.




Figure 9: Impact: full backups

The introduction of VM incremental backups in the NetBackup 6.5.2 release provides a number of advantages, including:

- Backup times are significantly reduced (see Figures 9 and 10)
- Backup windows are shortened
- The impact that backups have on the ESX server is reduced
- The amount of backup storage (disk/tape) needed is decreased




Figure 10: Impact: incremental backups

How does this work? With NetBackup 6.5.2 we combine two types of VMware Consolidated Backup backups. The first method involves copying over the complete vmdk files to the backup proxy. The second method simply creates and mounts a "shortcut" of the VM (no data is copied to the backup proxy during this mount operation). This shortcut represents the individual files that exist inside that VM. Single file backups are possible when defining and using this second method, but implementing this individual file backup method with simultaneous vmdk level backups is not possible.

The full backup uses the first, vmdk level backup. The incremental backup uses the second method. We can do this because we have already mapped and know exactly which individual files (for example, Word documents) exist inside the vmdk files that we designate as the full backup. For the incremental backup, we mount the shortcut of the VM, search for files that have changed since the full (vmdk) backup, and then backup *only* the files that have changed. The advantage of this technology is that for incremental backups, we copy to the backup proxy only the files that have changed since the vmdk full. This significantly reduces backup time, I/O load on the VMware datastore and amount of data written to tape or disk, while retaining the ability to restore individual files regardless of which backup method was used. This technology does not rely on any deduplication technology but can be used with any deduplication product.

There are some issues when using this methodology. Restoring to a specific point of time requires two separate restore processes (when using cumulative incrementals). The first restore is the restoration of

the full VM from the vmdk backup. The second restore is restoring individual files to the newly restored VM. Once both of these restores have been performed, all of the data associated with the VM-based on this restore point is safely restored. Incremental backups are file based (non vmdk) backups only. If a VM is patched or an application is installed or modified, we recommend performing a subsequent vmdk-based (full) backup as soon as possible to ensure that all system state information is safely captured.

Utilizing this incremental backup technology can dramatically increase backup performance while keeping all of your data safe and protected. Internal Symantec testing has indicated that incremental backups can be as much as 80% faster than comparable full backups.

It should be noted that incremental backups protect data, not the OS system state. When using this incremental backup technology it is recommended that periodic full (vmdk level) backups are scheduled and implemented, especially if OS level patches or any applications are installed or modified within the virtual machine.

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 website. 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 © 2009 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.

12/07 xxxxxxxx