



Confidence in a connected world.

Veritas NetBackup™ 6.5 for VMware 3.x Best Practices

Veritas NetBackup™ 6.5 for VMware 3.x

Best Practices

November 2007

Contents

Executive overview	4
Intended audience	4
Technology overview	4
Comparative technology overview	5
Three sample configurations	6
Configuration #1: Veritas NetBackup client installed inside the VMware Service Console	6
Configuration #2a: Veritas NetBackup client installed inside each virtual machine	7
Configuration #2b: PureDisk client installed inside each virtual machine	9
Configuration #3: Veritas NetBackup for VMware integrated with VMware Consolidated Backup ..	10
Implementation and deployment	11
Configuration #1: Veritas NetBackup client installed inside the VMware Service Console	11
Configuration #2a: Veritas NetBackup client installed inside each virtual machine	13
Configuration #2b: PureDisk client installed inside each virtual machine	14
Configuration #3: Veritas NetBackup for VMware integrated with VMware Consolidated Backup ..	15
Appendix A: Terminology	20
Appendix B: Additional resources	21
Appendix C: VMware Backup Proxy sizing	22

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 software is making its way into data centers of every size. Recognizing this trend, Veritas NetBackup™ has engineered innovative, award winning data protection solutions designed specifically for VMware environments. This paper covers best practices for designing solutions for—and protecting—VMware virtual machines.

As innovative as virtual machine technology is, it also introduces new data protection issues. Is it best to protect the virtual machine (VM) 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 relative advantages and disadvantages of these three backup configurations will be discussed in detail in this document.

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 virtual machines. Each of these technologies has relative advantages and disadvantages.

Technology overview

Two backup paradigms for VMware environments are discussed in this document:

On-host backup: These technologies involve installing a NetBackup or PureDisk Client inside each virtual machine or on the ESX Service Console. This backup methodology is popular because implementation is essentially the same as with physical machine backups. The downside is that backup activity on one virtual machine may also impact the performance of every virtual machine hosted on the ESX Server. This process is described in configurations #1, #2a, and #2b.

Off-host backups: This design takes advantage of the VMware Consolidated Backup (VCB) technology. Introduced with Virtual Infrastructure 3, this SAN or iSCSI based technology offloads backup processing from the ESX Server to a separate Backup Proxy server. NetBackup 6.5.1 adds Granular File Restore from full virtual machine (vmdk) backups, an award winning technology offered only by NetBackup. This technology is described in configuration #3.

Comparative technology overview

There are a number of ways to protect virtual machines. In this paper, we cover three of the most popular ways to protect 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.

Table 1. Solution comparison

	NetBackup client in Service Console	NetBackup client in virtual machine	PureDisk client in virtual machine	NetBackup for VMware VCB integration
Recommended for	DR (vmdk only) restores. Shared storage not available.	Single file and DB backups. Shared storage not available.	Single file with de-duplication advantages. Shared storage not available.	Low backup impact on highly impacted ESX servers in a shared storage environment.
VMDK level backups	✓			✓
Individual file backups		✓	✓	✓
LAN	✓	✓	✓	
SAN/iSCSI	✓	✓		✓
Cataloged & indexed Windows files		✓	✓	✓

Table 2. Performance comparison

	NetBackup client in Service Console	NetBackup client in virtual machine	PureDisk client in virtual machine	NetBackup for VMware VCB integration
Impact on virtual machine	●●●●●	●●●●●	●●●○○	●●○○○
Impact on ESX Server	●●●●●	●●●●●	●●●●○	●●●○○
Backup performance	●○○○○	●○○○○	●●●○○	●●●●●
Additional hardware requirements	●○○○○	●○○○○	●○○○○	●●●○○

Lower: ●○○○○

Higher: ●●●●●

Three sample configurations

Configuration #1: Veritas NetBackup client installed inside the 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 virtual machine (figure 1). Installing the NetBackup client inside the Service Console gives direct access to the files that make up the virtual machines (VMs)—the vmdk files. This method is easiest to implement if the virtual machines are powered off. In this state, the virtual machines are static and unchanging. If the VMs are powered on, additional pre-backup processing or scripting would be 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 (see Appendix B).




Figure 1. NetBackup client in Service Console

The relative advantages and disadvantages of this backup configuration are described in the following sections.

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.
- Entire VM restores are simple. The 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.

- The entire VM can be backed up by backing up the vmdk files.
- The configuration supports both LAN or 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 are likely if consistent backups of live VMs are required.
- Single (OS-level) file restores cannot be performed directly from the NetBackup interface. However, single file restores can be performed using VMware created tools such as “mountvm.exe.” See the resources in Appendix B for more information.
- This configuration does not support database or application backups.

Configuration #2a: Veritas NetBackup client installed inside each virtual machine

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 (see 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 like backing up a physical machine.




Figure 2. NetBackup client in virtual machine

Advantages:

- This is a simple and familiar implementation. Traditionally, most physical machine backups have been performed this way, which makes the transition to VM backups using this technology a straightforward 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:

- Entire virtual machine (at OS level) restores are more complex.
- FullVM (vmdk) restores are not possible.
- 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.
- Resource-intensive backups often place a heavy load on shared network and CPU resources.
- Client software installed on each VM needs to be maintained and updated.
- The VM must be powered on for backup processing to occur.

Configuration #2b: PureDisk client installed inside each virtual machine

This configuration is just like configuration #2a, except that a PureDisk client is used.




Figure 3. PureDisk client in virtual machine

Advantages:

- This is a simple and familiar implementation. Traditionally, most physical machine backups have been performed this way, which makes the transition to VM backups using this technology a straightforward task.
- Single file backups 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.
- Works with any storage type: SAN, NAS, DAS, or iSCSI.

Disadvantages:

- Entire virtual machine (at OS level) restores are more complex.
- Full VM (vmdk) restores 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. Backup scheduling should take this issue into account.
- Client software installed on each VM needs to be maintained and updated.
- The VM must be powered on for backup processing to occur.

Configuration #3: Veritas NetBackup for VMware integrated with VMware Consolidated Backup

NetBackup for VMware provides an alternate client backup technology for virtual machines with integrated file recover capabilities (see figure 4). Used in conjunction with VMware Consolidated Backup (VCB) off-host data protection technology from VMware, NetBackup 6.5 builds on and significantly enhances the VMware solution. Basic VCB integration is included with NetBackup 6.5. Enhanced Granular File Restore from a vmdk is included in NetBackup 6.5.1.




Figure 4. VMware Consolidated Backup

Advantages:

- Two restore options from a single backup: single file restore or entire (vmdk level) VM restore.
- Integration with the VMware Converter technology greatly simplifies restores of VMs.
- 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.
- There is no need to install NetBackup software on the VM or inside the ESX Service Console.
- Integration with VMware Virtual Center provides easy virtual machine discovery.
- The configuration is Distributed Resource Scheduling (DRS) and VMotion aware.

Disadvantages:

- The configuration requires a shared SAN or iSCSI environment and a separate W2K3 Proxy Server.
- Live (hot) application or database backups require additional pre- and post-backup processing to ensure backup data consistency.
- Granular file restore is currently limited to support for Windows® VMs.

Implementation and deployment

Configuration #1: Veritas NetBackup client installed inside the VMware Service Console

Installation procedure

There are two major components to installing this configuration:

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

The second task is technically not required—but it is recommended. It 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, but—because the backup is considered crash-consistent—restores are not assured. 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 the purpose.

Configuration

This deployment requires no specific hardware configuration. It does, however, have 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.

Configuring a NetBackup Policy: NetBackup policy configuration is reasonably straightforward, but a few issues should be kept in mind:

When defining backup selections, do not use the base VMFS directory on the ESX Server as a backup selection; instead, select a child folder within the vmfs directory. For example, to back up virtual machines that exist within the “datastore1” datastore, you would select the following path on the ESX Server as the backup path within the “Backup Selections” tab on the NetBackup policy definition:

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

Restoration procedure

Using this procedure, the VM vmdk files have been safely backed up. Restoring the VM based on these vmdk files is a two-step process. The basic procedure 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 using the vcbRestore command. The entire procedure is outlined in the “Virtual Machine Backup Guide.” See Appendix B for more information.

Hints, tips and best practices

There are several 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 sharing that datastore. This can be configured in the policy definition using the “Limit Jobs Per Policy” attribute (see figure 5).

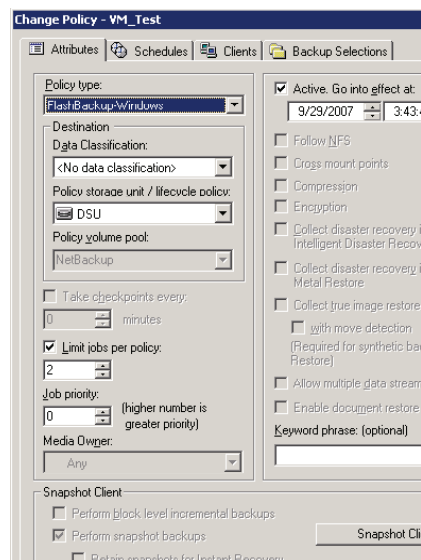


Figure 5. NetBackup policy configuration

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

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

Configuration #2a: Veritas NetBackup client installed inside each virtual machine

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 it would be for the OS hosted on a physical (not virtual) machine.

Configuration

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

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.

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. Accordingly, 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 includes Synthetic Backup technology that can be used to minimize or eliminate full backups. If only incremental backups are performed, only the data that has changed since the previous backup is copied to the NetBackup Media Server, which significantly decreases 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 providing incremental 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 sharing 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 virtual machine defined in a policy resides on the same datastore.

Configuration #2b: PureDisk client installed inside each virtual machine

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 it would be for the OS hosted on a physical (not virtual) machine.

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.

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 VM's targeted for backup, and the datastore or datastores that the VMs reside on.

Restoration procedure

One advantage of this configuration is that restores can easily be—and typically are—directed toward the original VM. The previously installed PureDisk 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. Backup policies should be defined to limit the number of simultaneous backup jobs that are running on each physical ESX Server.
- PureDisk de-duplication 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—which will, in turn, 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 virtual machine defined in a dataselection or policy resides on the same datastore.
- Limit the amount of dataselections and policies as much as possible.

Configuration #3: Veritas NetBackup for VMware integrated with VMware Consolidated Backup

Installation procedure

Configuration and installation of NetBackup for VMware are fairly straightforward. The following are the installation steps required to implement VCB within NetBackup (see figure 6):




Figure 6. VMware Consolidated Backup

1. Ensure that the hardware (especially the SAN) 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 VCB backups is a properly configured SAN environment. VMware SAN requirements can be found in the VMware SAN Configuration Guide, located at:
http://www.vmware.com/pdf/vi3_san_guide.pdf
2. On the Backup Proxy, install the supported version of the VMware Consolidated Backup Framework. The installation of this VMware component requires no specific configuration tasks.
3. Install the NetBackup component of choice (Master, Media, or Client) on the VMware Backup Proxy.

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

Configuration

This backup method entails two primary configuration tasks: configuration of the two VMware Consolidated Backup components within NetBackup. Those components are:

- **The Virtual Center Server**—or optionally the individual ESX Servers, if no Virtual Center Server is used
- **VMware Backup Proxy**—the component where the VMware images are mounted and backed

Both of these entities are defined within the NetBackup Admin Console. Consult the NetBackup for VMware Configuration Guide (see Appendix B).

Configuring a NetBackup policy: A NetBackup for VMware policy can be created either manually or with the help of the NetBackup snapshot wizard. With NetBackup 6.5, integration with either the ESX Server or the Virtual Center Server provides simplified virtual machine discovery.

Several NetBackup for VMware policy attributes are specific to VMware backups. Optimal backup performance can be achieved if these attributes are correctly applied. The attributes and their recommended settings are as follows:

- **Client definitions**—This policy attribute is probably the single most important policy setting in VMware environments. When clients are added to a policy, it is recommended that they 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 the associated datastore. This technique would allow for direct control over how many simultaneous backups

and subsequent VCB 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 (see 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 limits that amount of redo log creation—and all of this manages the I/O impact that backup operations have on the datastore.




Figure 7

- **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 kept 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 occur—or until the I/O impact on the VMs associated with the datastore begins to have a negative impact on the performance of the associated VMs.

Restoration procedure

Single file restores—NetBackup for VMware does not require that a NetBackup client be installed inside either the VM or inside 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 may be accessed and transferred to the VM through this share.

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 command `vcbRestore.exe`. Automatic VM restores can be performed by directing FullVM restores to a VMware Converter Server. The Converter Server automatically restores the VM to the specified ESX Server and registers the VM.

Hints, tips, and best practices

The success of VCB snapshot creation and deletion can be directly attributed to two things:

- The amount of I/O occurring on the VM datastore during snapshot creation
- The correct design of the I/O substructure associated with each datastore

To avoid snapshot-associated issues, backups should be scheduled during times of relatively low I/O activity on the VM. Reducing the number of simultaneous backups (and, in turn, VCB 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 documentation listed in Appendix B.

Group VMs in a policy with a single datastore. In other words, align NetBackup policies with each datastore. This allows you to manually isolate and control the amount of backup-related I/O that occurs per datastore—which, in turn, limits the impact that backup I/O has on the target VMs.

The VMware Backup Proxy can be configured as a NetBackup Master, Media, or Client. We recommend that the Backup Proxy be configured as a media server. The Backup Proxy is a natural

focal point for backup-related I/O. Backup Proxy access to VM files is typically made through a fast Fibre or iSCSI connection. In most configurations, it makes the most sense to avoid using only a NetBackup client to send VM backup data through the network—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 (see Appendix C), another Backup Proxy can be added to increase backup throughput capacity.

Once a VCB 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 is as fast as possible. The snapshot mount point should be configured over as many dedicated spindles as possible. 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-FullVM -t fullvm
```

Where the following VCB components are defined:

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-FullVM`

Multiple mount points (used to stage VCB snapshots) can be defined and used on the Backup Proxy server (see figure 8). Properly configured, additional VCB mount points 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.

Each mount point (used to stage VCB 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, promoting

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.

Upgrade to the latest version of VMware Virtual Infrastructure. This includes the latest version of ESX Server, Virtual Center Server, and VCB Framework. Newer versions of Virtual Center components typically have enhancements that improve VCB snapshot reliability.

Appendix A. Terminology

Backup Proxy—System designated as the off-host backup system. At a minimum, the Backup Proxy needs to have the VCB 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 virtual machine.

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

RDM—See Raw Device Mapping.

Sync Driver—Flushes OS buffers (Windows only) before VCB snapshots are initiated. The Sync Driver is installed using VMware Tools.

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.

vLUN driver—A VMware driver that translates vmdk files and represents them as individual OS-level files on the Backup Proxy. The vLUN driver exists on the Backup Proxy.

VM—An acronym standing for virtual machine.

VMDK—A designation specific to the files that comprise a VMware virtual machine. 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 virtual machine, VMware Tools enhance virtual machine performance and add additional backup-related functionality.

Appendix B. Additional resources

Storage/SAN Compatibility Guide For ESX Server 3.x. A critical component for making sure that the SAN environment is configured with supported equipment; it will help ensure that VCB style backups are as reliable as possible:

http://www.vmware.com/pdf/vi3_san_guide.pdf

I/O Compatibility Guide For ESX Server 3.x. Provides specific hardware and driver compatibility information for ESX 3.x:

http://www.vmware.com/pdf/vi3_io_guide.pdf

SAN Configuration Guide. Provides basic configuration information related to Fibre environments:

http://www.vmware.com/pdf/vi3_30_20_san_cfg.pdf

Consolidated Backup in VMware Infrastructure 3. An introduction to the VCB technology:

http://www.vmware.com/pdf/vi3_consolidated_backup.pdf

Virtual Machine Backup Guide. A VMware created paper that discusses several facets of virtual machine backups including VCB:

http://www.vmware.com/pdf/vi3_301_201_vm_backup.pdf

Veritas NetBackup Backup Planning and Performance Tuning Guide. Provides significant detail related to NetBackup Media Server (and, in turn, the Backup Proxy):

<http://support.veritas.com/docs/281842>

Veritas NetBackup for VMware Configuration. Describes how to configure NetBackup for VMware virtual machine backups with or without a Virtual Infrastructure environment:

<http://entsupport.symantec.com/docs/289771>

Appendix C. VMware Backup Proxy sizing

During VCB 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 virtual machine environment. Accurately characterizing the capacity of the Backup Proxy can be broken down into three main areas:

1. VCB data path—This is the entire data path that the VCB-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 VCB staging area for a FullVM style backup. This occurs as follows (see figure 7):

A. After the snapshot creation, virtual machine data flow occurs between the ESX datastore and the Backup Proxy. Both iSCSI and Fibre Channel technologies are supported for this connection.

B. Data is then transferred from the Backup Proxy to a disk staging area where backup processing occurs.

Keep in mind that both data paths (A and B) 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) has been completed.

2. Veritas NetBackup data path—With VCB 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 VCB staging area to the NetBackup storage unit should be a design concern.

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, Media, or 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 (see figure 8) and not first sent over the network—typically a relatively slower transfer medium than storage directly attached to the NetBackup Media server.




Figure 8

3. Backup Proxy sizing—The VCB 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/second 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.

More NetBackup-specific sizing information can be found in the NetBackup Backup Planning and Performance Tuning Guide that can be found here:

<http://support.veritas.com/docs/281842>

The VCB 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 backed up simultaneously

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

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

Keep in mind that the VCB snapshot area is only a temporary staging area and is reused as VM backups are processed. While an extremely small staging area can create a backup bottleneck in the VCB 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 VCB snapshot or backup data at a nominal transfer rate of 140 MB/second. Additional sizing information can be found in the Veritas NetBackup Backup Planning and Performance Tuning Guide (see Appendix B).

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

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 © 2007 Symantec Corporation. All rights reserved. Symantec, the Symantec logo, Veritas, and NetBackup are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Windows is a registered trademark of Microsoft Corporation in the United States and other countries. Red Hat is a registered trademark of Red Hat, Inc. in the U.S. and other countries. Other names may be trademarks of their respective owners.
011/07 13540287-1