WHITE PAPER:
NETBACKUP BEST PRACTICES

Symantec NetBackup™ Best Practices

NetBackup 7.1 for VMware

This document looks at and compares the various ways of protecting VMware® virtual machines (VMs) with NetBackup and discusses the advantages and disadvantages of each approach.

George Winter
Technical Product Manager

CONTENTS

1.0 Executive overview

In the past few years, much has changed relative to virtual-machine protection. Deduplication technologies are now an important part of protecting VMs. With embedded deduplication, Symantec NetBackup 7.1 makes implementing this ultra-efficient backup technology easier and more flexible than ever before. VMware vSphere™ 4 and the vStorage API for Data Protection (VADP) have significantly changed how VMware-based VMs can be protected. No other hypervisor vendor provides this advanced level of protection and no other backup vendor outside of Symantec provides better backup integration with this application programming interface (API).

Beginning with the 6.5 release, NetBackup has provided an advanced, award-winning backup technology designed specifically for VMware VM protection. A lot has changed since that release. Besides featuring a unique Granular Restore Technology™ for Windows® and Linux® VMs, NetBackup 7.1 also now features:

- Native support for VMware VADP
- Extremely fast and efficient block level incremental (BLI) backups
- Searchable, instant granular (single) file restores from both Windows and Linux VMs
- File-level restores directly from BLI-based backups
- Embedded “VMware aware” target- or source-based deduplication

Advanced VM protection technologies are certainly desirable but they can also introduce new data protection questions and challenges. Is it best to protect the VM by backing it up via a NetBackup deduplication client or using NetBackup for VMware VADP integration? The relative advantages and disadvantages of these backup configurations will be discussed in detail in this document. We will also provide guidance for successful implementation of these tools.

2.0 Backup technology overview

Virtual machine technologies provide a significant amount of flexibiltiy and scalability in the data center. Beyond these inherent capabilities, hypervisor-based systems also open new options for backup protection. Protecting an entire VM can be as simple as backing up a few Virtual Machine Disk (VMDK) files, but this level of simplicity comes at a certain cost. All-or-nothing VM restores don’t address every restore scenario. The typical backup administrator does not just need the ability to restore everything but they also need to be able to find and restore anything. Single-file restores are typically the most common type of restore request, but searching for and extracting individual files (Word doc, etc.) from a backed-up VMware VMDK file can be a time-consuming and resource-intensive activity.

Restore options should not be the only consideration when selecting backup methods. The backup method chosen has implications related to the backup impact on the host VMware ESX(i)™ system, the VMware ESX® datastore (storage) as well as the VM itself. In this paper, we will discuss two basic procedures that we suggest be deployed for VM protection. Each of these methods offers performance advantages as well as limitations. We recommend that each method be thoroughly examined so that the proper technology be selected for an acceptable balance between optimal backup performance and restore features.

Local backup: This involves installing a standard NetBackup client inside each VM and backing up the virutal machine the same way that would be used if it was a physical system. This backup methodology is popular because the implementation process is essentially the same as with physical machine backups.

Off-host backup: This method takes advantage of the VADP. Introduced with VMware vSphere 4, this method off-loads backup processing from the ESX server to a separate backup server (NetBackup refers to this system as the "VMware Backup Host"). NetBackup adds Granular Restore Techonology—an award-winning technology only offered by NetBackup—from both full and incremental VM (VMDK) backups. The vStorage API provides off-host backups with shared storage configurations using either Internet Small Computer System Interface (iSCSI) or storage area network (SAN) technologies.

2.1 NetBackup deduplication technologies

NetBackup 7.1 offers an unparalleled number of deduplication options, including embedded deduplication. These technologies include:

Media Server Deduplication Pool (MSDP): Embedded within the NetBackup 7.1 media server, MSDP offers both source- (client) and target- (backup destination) based deduplication. Both can be implemented within the standard NetBackup 7.1 client.

PureDisk™ Pool: This is an implementation of the classic PureDisk deduplication technology. PureDisk can be implemented as a stand-alone backup application with its own client.

NetBackup 5000 Deduplication Appliance: The NetBackup 5000 series of servers are stand-alone, preconfigured appliances that support both source- and target-based deduplication.

Each of these technologies can be deployed effectively in VMware VM environments. Sample deployments are described in the following sections.

3.0 Backup method comparison

Utilizing the two basic backup catagories described in the "Backup technology overview" section (page 1), we will discuss three suggested backup configurations for performing VMware VM backups. Table 1 on page 7 provides a comparision of the relative merits of each of these backup methods.

3.1 NetBackup (source) deduplication client installed in VM

In virtual environments, a traditional streaming backup with a standard client installed in the VM puts a tremendous strain on the VM host (ESX server). This is because the traditional backup process consumes significant central processing unit (CPU) and network bandwidth within the virtual infrastructure. NetBackup deduplication technologies address these challenges by offering a disk-based backup solution that dramatically reduces the size of backups and the network bandwidth required to perform them.

Providing more efficient client backups is accomplished by using a data deduplication technology at the source (client), which in this case would be within the VM guest OS. NetBackup deduplication identifies redundant sub-file data segments and only sends unique segments to the backup system. As a result, the NetBackup deduplication client eliminates the need to send or store duplicate files, such as OS files and executables, as well as redundant segments of files or database components as they are modified or changed over time. NetBackup deduplication also supports all major operating systems and databases, such as Microsoft® Exchange® and SQL Server®.

This architecture essentially gives you a backup solution that closely resembles that of physical servers, is simple to configure, and provides the same file-level restore capabilities of traditional backups. Running a NetBackup deduplication client inside the VM is a fully supported backup method. Standard OS support rules apply.

Advantages

- Simple and familiar implementation. Physical machine backups have been performed using client-based technologies for decades. Backup administrators find making the transition to VM backups using client-based technologies a straightforward task.
- Single-file or folder backups and restores for all supported guest OSs.
- Single-file restores directly into the guest OS are supported.
- Application backups (e.g., SQL Server and Exchange) are supported. The configuration process is exactly the same as configuring the same type of backup within a physical system.
- Extremely light, efficient data transfer and storage. Using deduplication technologies, storage for backups can often be reduced by 10–50 times. Bandwidth consumption for daily full backups is reduced by up to 500 times.
- Lower overall backup-induced resource utilization on the ESX host.
- Fully supports VMs configured using VMware Fault Tolerance (FT) technology.

Disadvantages

- Full VM images (VMDK files) are not backed up, making entire VM restores (i.e., disaster recovery) more complex.
- The backup processing load on one VM may negatively impact ESX resources available to other VMs hosted on the same physical server. It is recommended that backup scheduling be planned so that an excessive number of simultaneous backups never occur on a single ESX host.
- Client software installed inside each VM needs to be maintained and updated.
- The VM must be powered on for backup processing to occur.

3.2 NetBackup for VMware: Integration with the vStorage API for Data Protection

The VMware vSphere 4 release is a significant departure from the previous Virtual Infrastructure 3 (VI3) technology. There is an abundance of new features in vSphere 4, many related to VM protection. These improvements are delivered through the vStorage API for Data Protection. VMware has developed this API specifically for VM backups. Since the 7.1 release, NetBackup has provided full (native) integration of all the backup features that vSphere 4 and the VADP provide.

The VADP has additional architectural advantages as follows:

Off-host backup processing: Supported in shared storage (Fibre or iSCSI) environments.

No staging area required: NetBackup 7.1 no longer requires any disk for a staging area or holding tank. With NetBackup 7.1, the backup data stream is direct from the source ESX datastore to any destination storage unit type that NetBackup supports, including disk, tape, virtual tape library (VTL), or deduplication target (including the new NetBackup MSDP). By eliminating the requirement for a staging area (or holding tank), NetBackup 7.1 provides a significantly improved backup performance capability. The time-consuming step of temporarily staging backed up data to a disk staging area is no longer required. This applies to both ESX 3.2 U2 (and later) and ESX 4 (vSphere 4) systems.

Enhanced incremental backup technology: Another powerful feature of VADP is related to incremental backups. NetBackup 7.1 features a much faster and efficient BLI backup technology based on the VMware Changed Block Tracking (CBT) feature. A VM running on an ESX 4 server can be completely protected by backing up only the blocks that have changed since previous incremental or full backups. Benchmarking tests have indicated that NetBackup 7.1 BLI backup is extremely fast and efficient.

Better single file restores: NetBackup 7.1 fully supports restoring individual files (Windows and Linux VMs) directly from both incremental and full VMDK backups without the need to temporarily reconstruct the entire VM in order to gain access to individual files inside that VM.

While all vendors will be able to utilize the VMware CBT technology, only NetBackup is able to restore individual files or the entire VM directly from both full or BLI backups written to any backup destination. One hundred percent of this data is indexed and any file or VM can be searched for and instantly found and restored without having to first restore or restage the entire VM to disk.

Advantages

- Supports any ESX datastore storage configuration including network-attached storage (NAS), direct-attached storage (DAS), Fibre, or iSCSI.
- Provides off-host backup capabilities with Fibre and iSCSI-based shared storage.
- Disaster recovery (entire VM restore) is simple.
- With off-host (shared storage) backups, backup processing is off-loaded to the NetBackup VMware Backup Host.
- No client software needs to be installed on any VMware component (including the VM) to enable backup processing.
- One hundred percent of all file permissions are accurately retained during single-file restores (Windows and Linux).
- Supports advanced target-based deduplication.
- VMware Distributed Resource Scheduler (DRS) and VMware vMotion™ aware.
- Automatic VM backup selection/inclusion via Virtual Machine Intelligent Policy (VIP) (NetBackup 7.1).

Disadvantages

- Supports ESX 3.5 U2 or above.
- Virtual machine must be hardware version 7.1 if CBT support is desired (note: hardware version 4 VMs are also supported with the VADP but the CBT feature is not supported on hardware version 4 VMs).
- A VMware FT environment is not supported with the vStorage API style of backup.

3.3 Combine application client and VADP backups

Up to this point we have discussed using either standard client technologies or the VADP to protect VMs. But what about protecting VMs that are running complex applications that require special backup processing designed specifically for that application? Flexible, quick, and robust restore options are essential for ensuring availability of these important VMs. For these instances, a combination of both standard client and VADP technologies could be used.

In this implementation, NetBackup client and application (database) agents are configured to backup the OS and application. Both full and incremental backups are performed. Periodic VADP backups could also be performed to augment the client backups. The advantage of using the VADP is that the VM can be easily and quickly restored at the VMDK level. This greatly speeds up and simplifies disaster recovery restores and eliminates the time-consuming task of restoring the complex applications from a client-based backup.

This backup model would be typically used for a smaller subset of the VM environment. Virtual machines running important or mission-critical applications could benefit from using this technique. Other noncritical VMs would primarily still use only one backup method.

Two NetBackup policies would be required: one for the VADP and a second for the application agent backup. The schedule used for these policies would depend on the desired restore SLA. A sample configuration might use the application client policy to back up the VM on a daily basis with the VADP policy defined to back up the VM once a week.

This configuration will not work for every application type. It is suggested that you contact the application vendor to verify that this methodology is appropriate for the environment.

Advantages

- Ultimate restore flexibility.
- Object-level restores are possible. Database or application components can be fully restored from the client/app agent backup.
- Disaster recovery restores are simple and fast. Restoring the entire VM restores the application and all its components as well as the operating system.
- The time-consuming process of reinstalling and reconfiguring the application would not required when restoring at the image (VMDK) level.

Disadvantages

- Database agents must be installed and maintained within the VM.
- Two backup runs must be scheduled.
- Additional back-end storage is required to accommodate both backup styles.

3.4 Target- (destination) based deduplication

With target-based deduplication, backup data is first sent to the deduplication (backup) target. Once the deduplication target recieves the data it is then deduplicated. Target-based deduplication has an advantage in that this deduplication is independent of the method used to transfer data to the deduplication target. Any client type can be used with this type of deduplication, including older version clients that do not support source-based deduplication. The data can be sent via a standard backup client. No modifications to the standard client are required. The data can also be sent to the deduplication destination via any technology or API that provides the ability to prepare and transfer backup data. For VMs, VMware has created the VADP designed exclusively for VMware-based VMs. The VADP can prepare a VM for backup and automatically transfer the VM data to any backup application that supports VADP. The VADP is ideally suited for target-based deduplication technologies such as NetBackup 7.1 MSDP or the NetBackup 5000 series of deduplication appliances.

Advantages

- No backup client software needs to be installed on any VMware component (vCenter server, ESX host, VM) when the VADP is utilized.
- Resource requirements for deduplication processing is off-loaded from any VMware component.
- Fully supports VMs configured using the VMware FT capability.
- Supports all major VM (hypervisor) technologies.
- Less back-end storage is required.

Disadvantages

- All backup data must be first transferred to the deduplication target before deduplication processing occurs.

3.5 Summary of backup options

Table 1 provides a summary of the key features and capabilities. Numbered items are scored out of 5: 0 means low or limited and 5 means high or extensive.

Selection Criteria	NetBackup Deduplication Client in Virtual Machine	NetBackup for VMware VADP Configuration Network (NBD) Transport	NetBackup for VMware VADP Configuration SAN/iSCSI Transport	Combined App Client and VADP Backups
NetBackup Client Installation Required	✓	×	×	×
Individual File Restore	✓	✓	✓	✓
Individual File Search (Indexed)	✓	✓	✓	✓
Optimized Synthetic Incremental	✓	×	×	×
File-Level Incremental	✓	✓	✓	✓
Block-Level Incremental Backups	×	✓	✓	✓
LAN-based backups	✓	✓	✓	✓
Off-host (SAN or iSCSI) Backups	×	✓	✓	✓
Backup Performance	1	3	5	4
Backup Configuration Complexity	0	0	1	2
Additional Hardware Requirements	0	0	0	2
ESX Server Resource Impact	4	3	1	2
Restore Options	3	4	4	5
Disaster Recovery Complexity	3	3	0	0

Table 1: NetBackup for VMware backup technology comparison

4.0 Implementation, deployment, and best practices

In this section, we will provide details related to the actual deployment of the backup solutions mentioned in the “Backup method comparison” section on page 2.

4.1 Configuration 1: NetBackup deduplication client installed in VM

NetBackup features a number of different deduplication technologies, including classic PureDisk deduplication, NetBackup appliance-based deduplication, and deduplication embedded in the NetBackup media server, or MSDP. All three of these technologies support client- and source-based deduplication.

4.1.1 Installation procedure

Installation of the NetBackup deduplication client is straightforward regardless of which deduplication technology is chosen. The process of installing NetBackup client software inside a VM is exactly the same process used for physical systems. In NetBackup 7.1, the deduplication capability is built into the NetBackup client. The PureDisk deduplication client is a separate client install.

4.1.2 Configuration

With NetBackup 7.1 MSDP, client deduplication is enabled through the NetBackup client properties found within the “Master Servers” host properties tab on the NetBackup console. Double-click on the master server entry in the panel on the right. This will bring up the “Master Server Properties” interface. Select the “Client Attributes” tab. From the “Deduplication > Location:” pull-down menu (lower right), select “Always use client-side deduplication” (Figure 1).

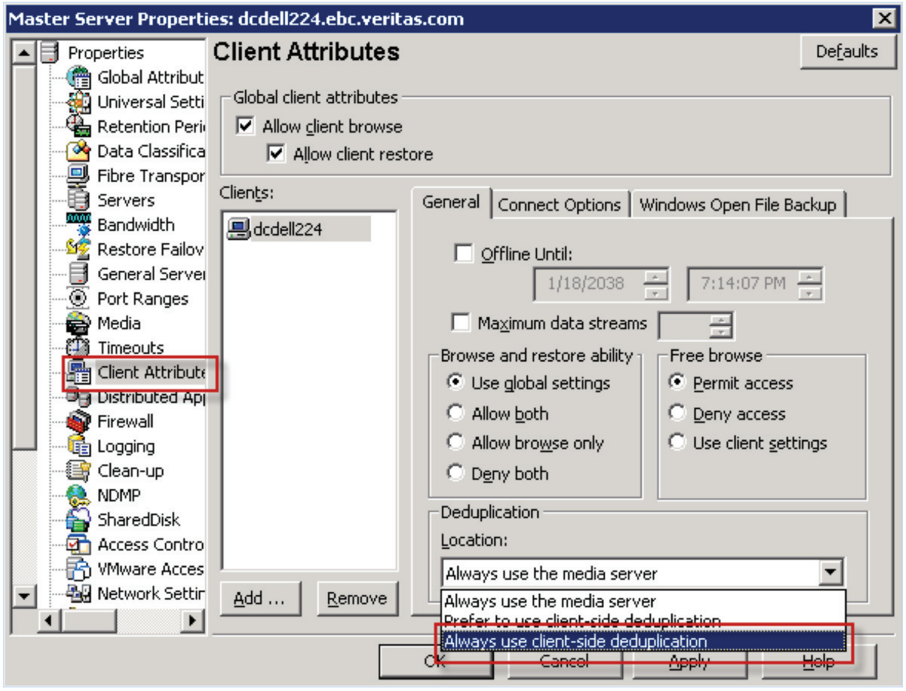

Master Server Properties: dcdell224.ebc.veritas.com

Client Attributes

Global client attributes

☑ Allow client browse

☑ Allow client restore

Clients: dcdell224

☐ Offline Until: 1/18/2038 7:14:07 PM

☐ Maximum data streams

Browse and restore ability: ● Use global settings ○ Allow both ○ Allow browse only ○ Deny both

Free browse: ● Permit access ○ Deny access ○ Use client settings

Deduplication

Location: Always use the media server

Always use the media server
Prefer to use client-side deduplication
Always use client-side deduplication

Figure 1: Configuring client-side deduplication

4.1.3 Configuring a client deduplication policy

Deduplication can be defined as either source- (client) or target-based deduplication. (Note: Target-based deduplication is covered in the following section.) Configuring the deduplication client is essentially the same regardless of whether it is deployed in a VM or a physical system.

The NetBackup MSDP deduplication option also supports optimized synthetic full backups. Optimized synthetic backups eliminate the need to perform periodic full backups. Full backup runs can negatively impact the VMs being protected. By using optimized synthetics, the resource impact that full backups create is eliminated. When backing up VMs using client technologies, optimized synthetics is highly recommend.

4.1.4 Restoration procedure

Restoring VMs from client-based backups is a very similar process as restoring physical systems from client-based backups. Both require specific OS knowledge. However, both can take advantage of the NetBackup Bare Metal Restore™ feature, which greatly speeds up and simplifies client-based restores.

4.1.5 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.
- 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.
- If possible, configure (align) NetBackup data selections and policies with the VMware datastores. To do this, create each policy so that every VM defined in a policy resides on the same datastore. See the “Balancing backup loads with the Virtual Machine Intelligent Policy technology” section (page 21) that describes the feature for additional information.
- Limit the number of simultaneous backups as much as possible. Try to achieve a balance between the number of active backups and achieving the backup performance necessary to protect your virtual environment while meeting your SLAs. Scheduling an excessive number of active backups can actually slow down overall backup performance.

4.2 Configuration 2: NetBackup for VMware integration with VADP

4.2.1 Installation procedure

One of the significant advantages of NetBackup vStorage API native integration is that configuring this technology is straightforward and simple to employ. The VADP is implemented as part of the NetBackup for VMware feature. NetBackup for VMware can be configured using a NetBackup master, media, or client system. Everything that is required to implement this technology is installed within any base NetBackup 7.1 install. No additional NetBackup or VMware software packages need to be added.

4.2.2 Configuration

There are only two steps that are required. First, the VMware Backup Host name (the NetBackup system that is designated for VM backups) is entered. Next the vCenter hostname and credentials are entered in the “credentials” section of the NetBackup master. These are the only configuration steps that are required. For any VADP-based backups, no NetBackup software needs to be installed on any VMware component, including the ESX/ESXi server, the vCenter server, or inside any VM. More information can be found in the “NetBackup for VMware Administration Guide” (see Appendix B).

4.2.3 Configuring a NetBackup for VMware policy

A NetBackup for VMware policy can be manually created or created via the “Snapshot Backup Policy” wizard. This section discusses some of the more important policy elements.

There are several NetBackup for VMware policy attributes that are specific to VMware backups. Optimal backup performance and reliability can be achieved if these attributes are correctly applied. Descriptions of these attributes are as follows:

Perform block level incremental backups: This is one of the most important attributes associated with vSphere 4 backups. Block level incremental backups provide the ability to perform true, incremental VM backups (Figure 2).

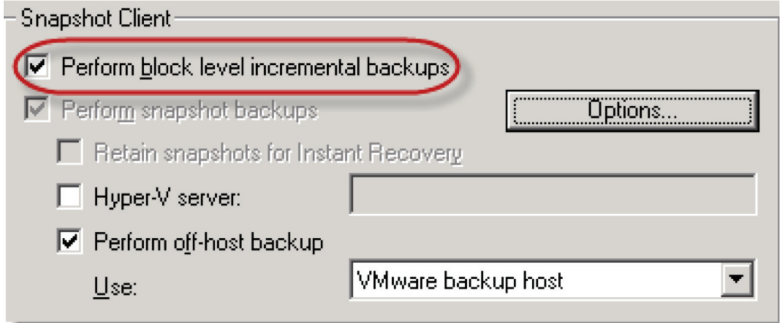
Snapshot Client

☑ Perform block level incremental backups

☑ Perform snapshot backups

□ Retain snapshots for Instant Recovery

□ Hyper-V server:

☑ Perform off-host backup

Use: VMware backup host

Figure 2: Block level incremental policy selection

The ability to perform true incremental VM backups is a significant backup technology enhancement. Previous VMware releases did not provide any method of performing incremental backups. These incremental backups utilize the vSphere 4 CBT feature. When enabled, BLI backups back up only the blocks that have changed since the previous full or differential/cumulative incremental backup. This technology provides significant advantages including:

- Incremental backups can be extremely fast. For detailed performance information, reference the "Symantec NetBackup™", Cisco® Unified Computing System (UCS), and VMware® vSphere™ Joint Backup Performance Benchmark" white paper (Appendix B).
- NetBackup is uniquely able to index, search for, and restore single files from BLI-style backups without the need to have direct access on disk to the previous full or incremental backups. This functionality is supported for Windows (NetBackup 7.1) and Linux (NetBackup 7.1) VMs.
- No restore functionality is lost when BLI backups is enabled. With NetBackup, any restore feature that is available for full backups is also available for BLI-style backups. Both single-file and entire VM restores can be processed from any full or differential/cumulative incremental style of backup. This technology is unique to NetBackup.
- Enabling BLI backups has extremely little impact on the VM. Blocks that have changed are kept track of in real time and do not have to be discovered before a BLI backup is performed. Very little disk space is required for the block tracking mechanism.
- NetBackup automatically enables the VMware CBT feature on the VM when BLI backup is selected within the NetBackup policy.
- BLI backups must be enabled before the first full backup of the VM is performed. If it is not enabled before the first full backup of the VM, subsequent incremental backups will not work properly.
- BLI-style backups are only supported on hardware version 7 VMs.
- BLI backup support is independent of the guest OS version.
- It is recommended that the BLIB functionality be enabled and utilized for all Windows or Linux (hardware version 7) VM backups.
- The VMware Backup Host must be configured to reference files based on "timestamp" instead of the "archvie bit." This setting only needs to be configured for the VMware Backup Host and does not need to be set for each VM.

Client Name Selection: The VM can be selected for backup based on three separate attributes: hostname (Domain Name System (DNS) name), vSphere display name (name displayed in the vSphere client), or Universally Unique Identifier (UUID).

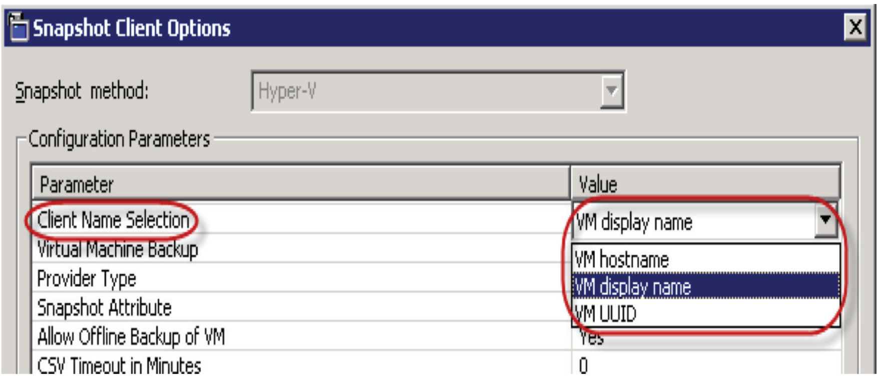

Snapshot method: Hyper-V

Parameter	Value
Client Name Selection	VM display name
Virtual Machine Backup	
Provider Type	
Snapshot Attribute	
Allow Offline Backup of VM	Yes
CSV Timeout in Minutes	0

Figure 3: Virtual machine Client Name Selection

There is no best choice for this backup attribute. Suggestions for their use are as follows:

VM hostname: This refers to the network (DNS) identification of the VM. When used, the hostname of the VM must be fully resolvable and correctly referenced in the DNS system. If the hostname is not correclty resolved, backups will fail. VMware tools must be installled on the VM so that the hostname can be registered to the vCenter server. If a VM has multiple Network Interface Cards (NICs) with different hostnames and IP addresses associated with each NIC, VM selection via the hostname can be problematic as the hostname reported by vCenter for the VM can be randomly selected from the available NICs configured on that VM. If this issue is encountered, it is recommended that either the VM display name or UUID be used for backup selection.

VM display name: This is the name that the VM is listed as within the vSphere client. The VM display name has no inherent relationship to the DNS hostname of the virutal machine. When the display name is used to reference the VM, it is recommended that the display name adhere to standard hostname naming conventions (i.e., no special characters in the name). NetBackup does support the use of spaces and some special characters in the display name.

Universally Unique Identifier: The UUID of a VM is a 32-digit alphanumeric designation for the VM. While the UUID can be complicated and not user friendly, in environmets where the VM hostname or display name are commonly changed or not accesible, the UUID can be the best way of referencing a VM for backup.

Regardless of which method is selected, keep in mind that if the name used to reference the VM is changed, subsequent backups will fail. If the name used to reference the VM is changed and modified in the NetBackup policy, subsequent backup data will be not be referenced to the original name of the VM possibly causing confusion during the restore process.

Virtual machine backup: This parameter defines NetBackup-specific technologies that can be applied to VM backups.

- Full VM backup: Standard backups of the VM. With this option, no special backup or restore functionality is enabled. There are no OS restrictions. Any guest OS supported by VMware can use this backup method. BLI-style backups are fully supported.
- Mapped full VM backup: Similar to the full VM backup option, selecting this parameter enables single-file indexing, searching, and restores from both Windows and Linux (NetBackup 7.1) VMs (Figure 3). All backup destination types (e.g., disk, tape, virtual tape library (VTL), deduplication) are supported for both backups as well as entire VM or single-file restores. No post backup processing is required and the VM data never needs to be restaged to disk as part of the restore proceedure. This option is highly recommended for both Windows and Linux VMs.

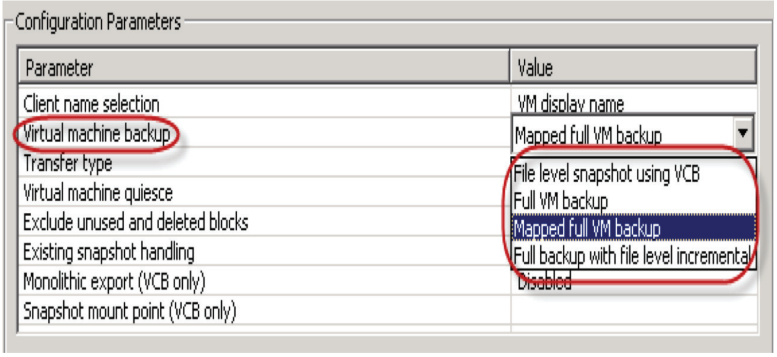

Configuration Parameters

Parameter	Value
Client name selection	VM display name
Virtual machine backup	Mapped full VM backup
Transfer type	
Virtual machine quiesce	
Exclude unused and deleted blocks	
Existing snapshot handling	
Monolithic export (VCB only)	Disabled
Snapshot mount point (VCB only)	

- File level snapshot using VCB
- Full VM backup
- Mapped full VM backup
- Full backup with file level incremental

Figure 4: Virtual machine mapping selection

Both the “File level snapshot using VCB” and “Full backup with file level incremental” options are specific to VMware Consolidated Backup- (VCB) style backups and are covered in the NetBackup 6.5 documentation. VMware has indicated that at some point in the future VCB will no longer be supported by the VADP. At that time, these two options will be deprecated.

Transfer type: This designates the transfer mode that will be used for the backup stream. The possible options include network-based defense (NBD) network, SAN (shared Fibre or iSCSI-based storage), or hot-add transfers. In most cases, the physical configuration of the storage will mandate the transfer type. We recommend that the desired transfer type be explicitly defined in the policy. For example, if the backup environment is configured for shared storage–based transfers then the SAN transfer type should be selected. The “try SAN then NBD” transfer mode should be used only in environments where the backup must succeed even if SAN transfers fail for any reason. When “try SAN then NBD” is selected, if the SAN transfer fails, the backup will attempt network-based transfers. This increases the chance that the backup run will succeed.

Virtual machine quiesce: At the beginning of backup processing on Windows VMs, VMware tools calls the Windows Volume Shadow Copy Services (VSS) provider to quiesce the VM. This ensures that the VM backup will be consistent and not corrupt. It is recommend that this option be enabled.

There can be conditions in which the VSS quiesce process can take a long time, timeout, or simply fail. This can occur for newer Windows releases, because of incompatibility issues with VSS, because specific applications are running inside the VM or the VM is extremely busy. If any of these cases cause the backup process to fail, the VM quiesce can be disabled through this option. This option can also be used to disable the VSS provider when troubleshooting snapshot issues associated with a problematic VM backup.

Exclude unused and deleted blocks: When any data within a VM is deleted, the blocks where the data resided are not deleted. The vStorage API backs up these blocks as part of the backup process. When enabled, this feature automatically excludes any of these unused blocks from the backup process. This capability can significantly increase backup efficiency and speed by skipping these unused and unecssary blocks. This feature is enabled by default and works with both thick- and thin-provisioned VM disks. It is recommended that this feature be enabled always. It should only be disabled when troubleshooting backup issues.

Existing snapshot handling: There are many VMware technologies that can create snapshots on a VM. The VADP creates a temporary snapshot as part of the backup process. Other VMware features such as DRS and vMotion can create VM snapshots as well. This option determines a course of action when an active backup-related snapshot is encountered at the beginning of the VM backup process (Figure 5).

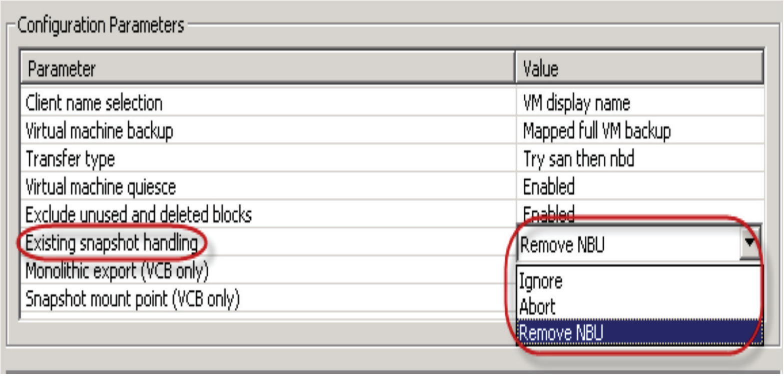

Configuration Parameters

Parameter	Value
Client name selection	VM display name
Virtual machine backup	Mapped full VM backup
Transfer type	Try san then nbd
Virtual machine quiesce	Enabled
Exclude unused and deleted blocks	Enabled
Existing snapshot handling	Remove NBU
Monolithic export (VCB only)	
Snapshot mount point (VCB only)	

Figure 5: Virtual machine snapshot handling

Three values are supported with this option (Figure 5):

Ignore: To ignore any existing backup-related snapshots and attempt to create a new snapshot.

Abort: If an existing snapshot is detected, the backup will abort. This would be used if there is the possibility that two NetBackup policies could be simultaneously backing up the same VM.

Remove NBU: When selected, this option can increase backup reliability. If a previous backup created a snapshot that was not properly removed when the backup finished, NetBackup will automatically remove this straggler snapshot at the beginning of the next scheduled backup of that VM. Selecting this option can increase backup reliability by eliminating issues associated with leftover snapshots. It is highly recommened that this option be selected.

The remaining two options, “Monolithic export” and “Snapshot mount point,” are specific to VCB-style backups and are not used when native vStorage API backups are implemented. VMware has indicated that at some point in the future VCB will no longer be supported by the VADP. At that time, these two options will be deprecated.

4.2.4 Restore procedures

NetBackup 7.1 provides greatly simplified restore offerings for VMware environments. Two basic restore options are available: individual file restores (for Windows and Linux OSs), and entire VM restores (any OS). These two restore options are described as follows:

Single-file restores: During the backup process, NetBackup for VMware does not require or in any way use NetBackup client software. If single-file restores directly into the VM are desired, a NetBackup client can optionally be installed without incurring any additional client license expense. Alternatively, if no client is installed inside the VM, an alternate client restore can be performed to a Windows share. The restored files are then transferred to the VM via access to this share. Even though the files to be restored are based on VM VMDK backups, the single-file restore interface and procedure is exactly the same as it would be if the VM had been backed up using traditional client-based backup technologies.

Full VM restores: NetBackup 7.1 provides a simple VM restore interface that walks the user through the entire restore process. The restore procedure is completely automated by NetBackup. No additional software needs to be installed anywhere and no additional hardware (e.g., staging disk) needs to be configured to implement entire VM restores. NetBackup automatically restores the VM to the specified ESX server and automatically registers the VM on the vCenter server. Once this fully automated process is complete, the VM is immediately ready for use.

Cloning VMs: An obvious use of full VM restores is to re-create the VM from any point in time due to a disaster-related event. But this restore technology can also be used for other purposes, such as cloning VMs. Cloning a backed up VM can be accomplished by selecting the “Alternate location” restore option from the “Virtual Machine Recovery” wizard. Simply append or rename the VM name in the “Display Name” field of the “Virtual Machine Recovery” wizard and select the original or an alternate vCenter/ESX(i) environment to restore to. The original VM will not be written over or damaged in any way. It is also recommended that VMs be restored to a different network (vSwitch) environment so that once powered on there are no IP or network conflicts.

4.2.5 Hints, tips, and best practices

- Internal testing and customer feedback has indicated that the success rate when using the VADP is markedly better when compared to the VMware CBT. It is highly recommended that you use the VADP method for most VM backups.
- Excessive Input/Output (I/O) contention can always be a cause of failed snapshots, and in turn, failed backups. To help limit the amount of backup-related I/O, it is recommended that the number of VM backups that are active per ESX datastore is limited. For example, if 20 simultaneous backups are desired across VMs distributed across five ESX datastores, backup performance and reliability would be improved if no more than four backups per datastore were configured. This would still achieve the goal of 20 simultaneous backups but would help equalize the backup-related I/O load across all ESX datastores.

- It is important to define NetBackup policy attributes such as “limit jobs per policy” so that the number of backup jobs run against each ESX(i) datastore is not excessive and in turn does not create unnecessary backup-related I/O contention. However, being able to accurately do this in dynamically changing real-world conditions can be a daunting task. The VIP feature introduced in NetBackup 7.1 can be of great help in this area. VIP can dynamically adjust backups so that no single ESX component is overtaxed by backup-related I/O.
- The VMware Backup Host can be configured as a NetBackup master server, media server, or enterprise client. As a general rule, we recommend that the VMware Backup Host be configured as a media server but all three configurations are fully supported.
- MSDP can be an extremely effective tool for reducing backup storage requirements and increasing overall backup speeds. We highly recommend the use of MSDP or other NetBackup deduplication technologies for improving overall backup efficiencies. The “Symantec NetBackup, Cisco UCS, and VMware vSphere Joint Backup Performance Benchmark” white paper discusses the performance advantages that can be achieved with MSDP.
- Separate host bus adapters (HBAs) should be used for each I/O path on the VMware Backup Host. For example, the destination storage unit (disk/tape) and the connection to the datastore(s) should be configured on separate HBAs. This ensures that there is no I/O contention that would be encountered if a single HBA was used. It is also recommended that each HBA be located (if possible) on separate internal buses within the VMware Backup Host server.
- If possible, upgrade to the latest version of vSphere. This includes the latest version of ESX server and the vCenter Server. Newer versions of vSphere components typically have enhancements that improve overall performance and snapshot reliability. Newer versions of ESX servers typically provide updated VMware Tools and VSS providers. Both of these upgrades can contribute to more reliable backups.
- BLI backups are your friend. NetBackup benchmarks have established that the vSphere CBT feature (implemented by NetBackup as BLI backup) can provide an extremely fast and efficient method of incrementally protecting VMs. Because NetBackup 7.1 enhanced granular file restore technology has been extended to include BLI backups on Windows and Linux VMs, there is no loss of restore functionality. Both single-file and entire VM restores can be processed from full as well as incremental backups. BLI backups are supported on any VM that is hardware version 7.
- Virtual machines configured using VMware FT technology cannot be snapshot and therefore cannot be protected using the VADP. Standard client-based backups are the best method for protecting these VMs.
- Be sure that the vCenter Server is protected. The vCenter Server can be configured within a VM or on a physical system. The VADP can be used to protect a vCenter Server configured in a VM but client-based technologies are best for vCenter systems configured on physical hosts.

4.3 Configuration 3: Combined application client and VADP backups

For VMs running important applications, individual file restores or entire VM restores may not provide an adequate level of restorability. For vital mission-critical systems, any and all restore options may be required. The restore options mandated by these important applications could include single-file restores, object- (DB/app) level restores, and entire VM disaster recovery (DR) restores. Combining both client and DB agent backups (for object-level restores) with vStorage API (for disaster recovery (DR)) VMDK backups can provide the ultimate in VM and application restorability. This type of VM protection is described in this section.

4.3.1 Installation procedure

Configuring the vStorage API VM backups within a NetBackup environment is extremely simple. The second part of this configuration involves installing and configuring the NetBackup application/DB agent. While agent installation occurs within a running VM, the installation procedure is exactly the same as if performed within a physical system. The appropriate application/DB agent documentation covers this in detail.

4.3.2 Configuration

As mentioned, configuring the application/DB agent is straightforward. What makes this method different is the need for proper backup scheduling. It is recommended that the vStorage API backup and client/DB agent backup be scheduled at different times to avoid resource conflicts within the VM that could be caused by simultaneous backups. With the knowledge that the vStorage API style of backup provides a simplified disaster recovery methodology and client backups provide excellent object-level restores, backups should be scheduled to take this into account in accordance with your SLAs.

4.3.3 Configuring NetBackup policies

Two separate NetBackup policies need to be configured. The vStorage API policy should be configured using the full VM backup selection within the policy “options” tab. This method does not perform individual file mapping and no single-file restores are possible, but it does provide for extremely efficient incremental (BLI) backups. It is the application or DB backup policy that will be responsible for protecting and restoring individual files. The DB/application policy would be configured as required by the application being protected. In the interest of reducing backup resource contention, care should be taken so that the execution schedule of each of these policies does not overlap.

4.3.4 Restoration procedure

Two very different and powerful restore options are available when deploying this backup method. For DR, restoring the entire VM from VADP backups is the best choice. A VM restore wizard makes this process easy and is part of the standard NetBackup product. Restoring the VM and the application/DB within that VM is a greatly simplified process. If restoring application/DB objects is necessary, the standard NetBackup application/DB restore interface can be used.

4.3.5 Hints, tips, and best practices

- Recovery Time Objective (RTO) is defined as the time required to restore from a data loss event. The method used to back up data has a direct impact on the RTO. Client-based application backups can provide for quick object-level restores but slower catastrophic (entire VM) restores are more complicated and take longer from this type of backup. Restores from a vStorage API backup are simple, but restoring the entire, large VM to recover a small, single application object can be overkill. RTO design decisions will vary, but understanding the RTO needs and balancing them with the correct backup strategy will ensure that recovery objectives are met.
- Be careful not to overdeploy this two-pronged approach. This backup method involves two backup runs. This translates into additional impact on the ESX system resources, longer backup windows, and more backup storage resources. Using the vStorage API style of backups will be more than sufficient for protecting and providing powerful restore options for the majority of VMs. More important, mission-critical VMs should be targeted for this style of backup.

5.0 Additional topics

5.1 VMware backup host sizing and performance

With the older VCB technology, sizing the VMware Consolidated Backup Proxy media server was a very different task than sizing the NetBackup 7.1 VMware Backup Host. With VCB backups, special care had to be taken when configuring the staging area (holding tank). During the backup process, backup data was temporarily staged to this holding tank and then eventually moved to the backup target. This meant that all backup data had to be moved twice during every backup run. Protecting 100 TBs of VM data required moving 200 TBs. A poorly configured staging area negatively impacted the performance of all VM backups.

The VADP eliminates the need for this staging area. This single design improvement has significantly enhanced the overall performance of VMware backups. It has also simplified the exercise of sizing the VMware Backup Host. Sizing the VMware Backup Host is now essentially the same process as sizing any NetBackup system. Detailed media server sizing information can be found in the “NetBackup Planning and Performance Guide” (see Appendix B). Standard media server configuration guidelines associated with CPU selection, I/O handling, and RAM sizing apply directly to the VMware Backup Host.

Benchmark testing using the NetBackup for VMware feature in conjunction with the VADP has indicated that both of these technologies scale extremely well. If performance issues are encountered, we suggest that you run basic performance tests to determine which hardware component is causing the performance issue. Virtual machine backup performance using NetBackup will be highly dependent on the available I/O bandwidth that is available both within the ESX environment as well as hardware and I/O resources available to the NetBackup VMware Backup Host.

Virtual machine backup I/O performance can be broken down into three specific areas: (see Figure 6)

1) Streaming data from the ESX datastore: Whether the backup stream occurs through the ESX server network or over shared storage (as detailed in Figure 5), backup data has to be read from the ESX datastore and then sent to the NetBackup VMware Backup Host.
2) Sending backup data to the VMware Backup Host: This connection can be either network-based or based on shared storage (e.g., Fibre or iSCSI). The speed of this connection determines how fast data can be ingested into the VMware Backup Host.
3) Backup target write performance: This target can be any backup destination that is supported by NetBackup, including disk, tape, VTL, or deduplication targets.

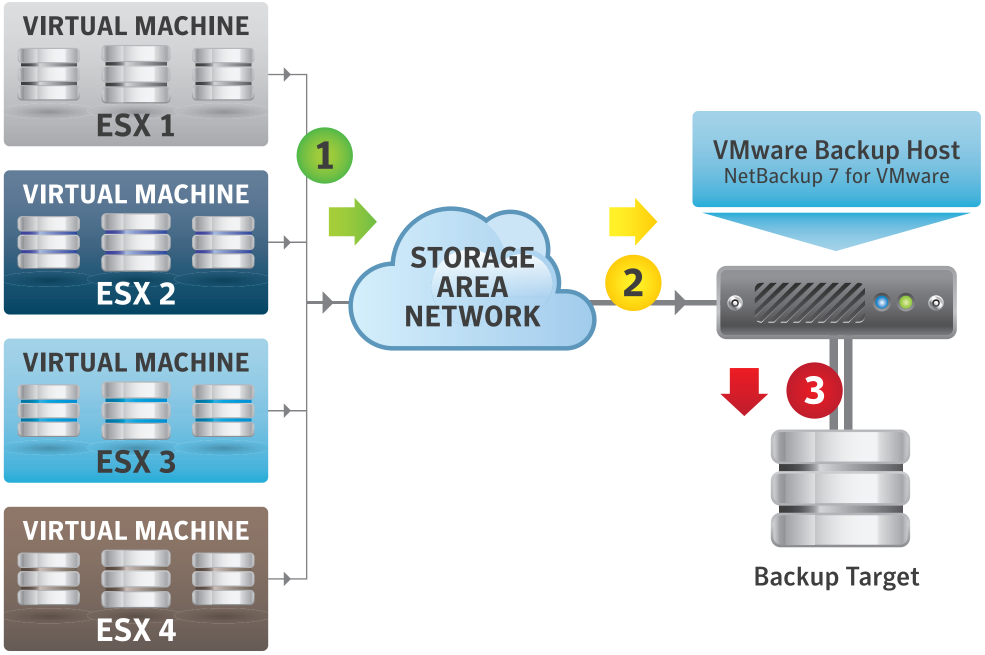

Figure 6: Virtual machine backup performance and sizing

The slowest performance of any of these three areas will determine the maximum performance that the NetBackup VMware backup host will be capable of. The "Symantec NetBackup, Cisco UCS, and VMware vSphere Joint Backup Performance" white paper (Appendix B) provides suggestions for testing and determining the performance of each of these three areas. We highly recommend referencing this paper for more detailed information. Optimzing the I/O capabilities of the backup environment will ensure faster backups, lessen the impact that backups have on the VMware environment, and reduce backup hardware requirements.

5.2 VMware Backup Host configured using a NetBackup client

While it is recommended that the VMware Backup Host be configured as a media server, there are several use cases in which this might not be optimal. For these instances, the VMware Backup Host can be configured as a NetBackup client or a NetBackup SAN client instead of a master or media server. In either of these cases, backup data is sent to the NetBackup client that then forwards it to the media server where it is written to any supported backup destination. This configuration is completely supported and in no way modifies the restore process or limits any backup or restore functionality.

The following is a list of configurations in which it might make sense to configure the VMware Backup Host as a NetBackup client:

Non-Windows media servers: Environments where media servers are exclusively non-Windows.

NetBackup appliances: As the NetBackup appliance OS is not Windows-based, it cannot currently be configured as a VMware Backup Host. In this situation, a NetBackup client could be configured as the VMware Backup Host. The client would simply forward all backup data to the NetBackup appliance without any loss of backup or restore functionality.

Media server failover/load balancing: Because the VMware Backup Host definition is static, media server failover/load balancing will not work when the media server is defined as the VMware Backup Host. For these situations, a NetBackup client can be defined as the VMware Backup Host. In this configuration, media server failover/load balancing will be fully supported.

SAN client: When the VMware Backup Host is defined as a client, the network connection between the VMware Backup Host client and the media server can be a performance bottleneck. Implementing a SAN client provides a direct SAN connection between the client (in this case a VMware Backup Host) and the media server. The fast SAN connection can eliminate the network as a bottleneck, improving overall backup performance. Backups can be processed at Fibre speeds.

5.3 Utilizing the hot-add backup method

NetBackup can be installed inside VMs with that VM configured as a VMware Backup Host (the result is essentially a backup appliance). This configuration is useful when additional physical hosts are not available or practical. The downside of using a VM for backups is that all of the backup I/O that is performed by this VM impacts the host ESX server and in turn impacts every VM hosted on that ESX server. Even though the target VM and the backup VM are hosted on the same ESX server, the backup stream must use the ESX host TCP/IP stack.

Enter the VMware hot-add feature. Hot-add is a transport mode that bypasses the ESX host's TCP/IP stack. With hot-add, the backup data path moves from the ESX datastore (storage) directly to the VMware Backup Host. This significantly improves backup performance while reducing the backup load on the ESX server. Hot-add supports any style of storage technology including DAS, NAS, Fibre, and iSCSI.

When possible, implementing the hot-add transport method is highly recommended when a VM is defined as the VMware Backup Host. It requires that the ESX server that hosts this VMware Backup Host has direct access to the ESX datastores that house the VMs that are targeted for backup. This means that the VM designated as the hot-add system can only back up VMs that reside on datastores visible to the host ESX server.

5.3.1 Hints, tips, and best practices

- When using the hot-add transport, there is no loss of backup and restore functionality.
- Hot-add is useful for remote offices where it is preferred to keep the backup data locally. Hot-add also eliminates the need to purchase additional backup system hardware.
- Keep in mind that VM-based backups impact the host ESX server. For a severely overtaxed ESX environment, physical host-based backups should be considered as an alternative to VM-based backups.

5.4 Balancing backup loads with the Virtual Machine Intelligent Policy technology

For decades, traditional client backups have been defined using the hostname of the client. This has worked well when the relationship to the physical host and the operating system was a one-to-one, direct relationship. Virtual machine technologies changed this physical one-to-one relationship. Dozens of operating systems can now reside on a single physical (ESX) host connected to a single storage LUN. Standard backup policy definitions do not translate well into this virtual environment. NetBackup engineering has designed the VIP feature to take into account the physical aspects of VM environments when designing backup policies.

VIP dynamically discovers VMware VMs and automatically places them in policies that are designed to take into account the physical attributes and performance limitations of the VM environment. The VIP technology limits and balances the backup impact so that no single VM or ESX server is unfairly impacted during backup processing.

VIP backups can be defined according to over 25 different VMware VM attributes. For most backups, we recommded that VIP backups be defined according to either the ESX server or according to the ESX datastore (storage).

When using SAN- (shared storage) based transfers we suggest that VIP backups be defined at the ESX datastore level. SAN backups pull data directly from the ESX datastore to the VMware Backup Host. When backups are defined per datastore, backup processing can be balanced across all datastores so that no single datastore is unfairly impacted by backup processing.

When using NBD or networked-based transfers, we recommend that VIP backups be defined at the ESX level. Since this network-based backup transfers data through each ESX host, this would have the effect of balancing network backup traffic across each ESX server. In this configuration, backup processing never unfairly impacts a single ESX server yet fast backup performance is still achieved.

APPENDIX A: Glossary

Backup proxy: A Windows-based system designated as the off-host backup system. With the release of NetBackup 7.1, the backup proxy designation is no longer used. See VMware Backup Host.

BLI backups: Block level incremental backups. This is the NetBackup 7.1 implementation of the Changed Block Tracking feature introduced with vSphere 4. When BLI backups is selected, NetBackup automatically enables the vSphere 4 Changed Block Tracking feature.

Changed Block Tracking: Introduced with vSphere 4, this ability provides true block level incremental backup technology for ESX 4 (HW version 7) VMs.

Datastore: The ESX server storage. Virtual machine (VMDK) files are created and stored on the ESX datastore.

Guest OS: The actual operating system that resides within the virtual hardware (i.e., the VM).

Holding tank: An NTFS formatted disk volume that is created on the backup proxy. This volume is used as part of the VMware Consolidated Backup process. As of the NetBackup 7.1 release, the holding tank is no longer required when vStorage API for Data Protection style backups is implemented.

Hot-add: A transport mode that enables a backup system configured as a VM to back up other VMs residing on storage visible to a common ESX Server. Using the hot-add transport can be faster than LAN backups as the backup data path is directly from the ESX datastore to the hot-add VM, bypassing the LAN and the ESX(i) TCP/IP stack.
Media Server Deduplication Pool: Introduced with the NetBackup 7.1 release, MSDP is a deduplication technology that is embedded in the NetBackup media server. MSDP features both source- (client) and target-based deduplication.

Staging area: See holding tank.

VCB: See VMware Consolidated Backup Framework.

VADP: See vStorage API for Data Protection.

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

Virtual Machine Intelligent Policy (VIP): An innovative VM selection technology introduced in the NetBackup 7.1 release. VIP can automatically select for backup newly introduced or moved (via DRS or vMotion) VMs. VIP can also load-balance VM backups, ensuring optimal backup performance.

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 Backup Host: The NetBackup system that is designated for backing up the VM environment. This Windows-based system can be configured as a NetBackup master server, media server, or client. VMware Backup Host is a designation specific to NetBackup.

VMware Consolidated Backup Framework: An off-host backup framework created by VMware. VMware Consolidated Backup has been superseded by the vSphere component VADP. The VMware Consolidated Backup Framework software package is no longer required.

VMware Tools: Installed inside each VM. VMware Tools enhances VM performance and adds additional backup-related functionality. See VSS Writer.

VSS Writer: VMware replaced the Sync Driver with a Volume Shadow Copy Service (VSS) writer beginning with the ESX 3.5 U2 release.

vStorage API for Data Protection (VADP): The VADP is a VMware technology available to backup vendors. Backup software vendors can integrate with it to perform centralized VM backups without the disruption and overhead of running backup tasks from inside each VM.

APPENDIX B: Additional Resources

vSphere 4 documentation: A generic link to ESX and ESXi documentation.
http://www.vmware.com/support/pubs/vs_pubs.html

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

VMware SAN Configuration Guide:
http://www.vmware.com/pdf/vsphere4/r41/vsp_41_san_cfg.pdf

VMware vStorage APIs for Data Protection: The vStorage APIs for Data Protection enable backup software to perform centralized VM backups without the disruption and overhead of running backup tasks from inside each VM.
http://www.vmware.com/products/vstorage-apis-for-data-protection/

VMware Virtualization Performance Resources: Learn more about VMware technologies designed to improve performance.
http://www.vmware.com/technical-resources/performance/

Overview of support for NetBackup 7.x in virtual environments: Details all aspects of the VM support that NetBackup provides.
http://www.symantec.com/docs/TECH127089

Symantec NetBackup 7.1 for VMware Guide: Administrator guide for NetBackup 7.1 VMware functionality.
http://www.symantec.com/docs/DOC3663

Symantec NetBackup, Cisco UCS, and VMware vSphere Joint Backup Performance Benchmark: Documents all aspects of this recent backup benchmark joint effort with Symantec, Cisco, and VMware.
http://www.symantec.com/business/netbackup

NetBackup Planning and Performance Guide: Detailed information related to sizing and tuning NetBackup systems for maximum performance.
http://www.symantec.com/docs/TECH62317

More information

About Symantec

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

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

For specific country offices and contact numbers,

please visit our website: www.symantec.com

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

Confidence in a connected world.