



OPENBENCH LABS



**Analysis:**

Better VMware vSphere Data Protection  
Through Unified Backup and Replication

# Analysis: Better VMware vSphere Data Protection Through Unified Backup and Replication

Author: Jack Fegreus, Ph.D.  
Managing Director  
openBench Labs  
<http://www.openBench.com>  
January 8, 2010

*Jack Fegreus is Managing Director of openBench Labs and consults through Ridgetop Research. He currently serves as CTO of Strategic Communications, Editorial and contributes to InfoStor and Virtualization Strategy Magazine. He has served as Editor in Chief of Open Magazine, Data Storage, BackOffice CTO, Client/Server Today, and Digital Review. Previously Jack served as a consultant to Demax Software and was IT Director at Riley Stoker Corp. Jack holds a Ph.D. in Mathematics and worked on the application of computers to symbolic logic.*

# Table of Contents

Executive Summary	04
Data Protection Test Bed	06
Backup & Replication Performance	09
Data Protection Process Value	17

# Executive Summary

“Veeam Backup & Replication simplifies sophisticated IT backup and recovery procedures in virtual environments that include both ESX and ESXi hosts—a capability that is precisely what IT needs to meet Service Level Agreements (SLAs) with aggressive RTO and RPO requirements.”

To establish highly efficient operations, the consensus among CIOs is that virtualization of systems and storage is the magic bullet for achieving higher IT resource utilization and lowering management costs. That has administrators at many IT sites now working with limited numbers of abstract device pools rather than multiple instances of proprietary devices. Nonetheless, system and storage virtualization also introduces multiple levels of logical abstraction and resource redirection that can obscure and complicate important IT operations.

## openBench Labs Test Briefing: Veeam Backup & Replication 4.1 Software

- 1) Minimize vSphere disk risk with backup and replication:** Veeam Backup & Replication combines backup and near-CDP-level replication in one interface to unify both data protection processes for IT with a synthetic backup process and consistent process wizards.  
**Backup Benchmarks: A full backup of a VM running Windows Server 2003 averaged 189 MB per second and an incremental backup of the same VM averaged 245 MB per second on a VMware® ESX 4.0 host.**
- 2) Maximize backup performance with deduplication and compression:** Data deduplication, compression, and recognition of thin provisioning lowered the footprint of a full backup of 8 VMs by a factor of 9 to 1.
- 3) Full support for ESX and ESXi without VCB:** Veeam Backup & Replication can leverage the new vStorage API or VMware Consolidated Backup (VCB) when performing backup or replication, allowing SMB sites with or without shared storage to fully leverage ESXi, including the ability to replicate VMs to an ESXi host.
- 4) Full support for VMware vSphere Changed-Block Tracking and thin provisioning:** Veeam Backup & Replication dramatically increases full and incremental backup speed, as well as reduces storage requirements.
- 5) Higher RPO support via Virtual CDP with replication:** Via support for synthetic backup, Veeam Backup & Recovery can leverage inline data deduplication and compression to support very narrow backup windows, which provides for near continuous data protection (CDP) with replication.
- 6) Ensure transaction consistency at VM and application levels:** Veeam Backup & Recovery employs VMware snapshots to ensure VM transaction consistency and Volume Shadow Services (VSS) for VSS-aware Windows applications such as Active Directory, SQL Server, and Exchange.

Among the hardest hit IT operations are those associated with file-level data protection. That puts compliance with regulatory mandates to secure and maintain critical business data at a granular file level in the cross hairs of IT's magic bullet for gaining operating efficiency. This situation complicates two important concepts that IT must balance when making the business impact analysis required for any IT disaster recovery plan: the Recovery Time Objective (RTO)—the maximum period of time that it could take to recover—and the Recovery Point Objective (RPO)—the maximum amount of data measured in time prior to the disruption that could be lost in the recovery process.

Veeam Backup & Replication software is specifically designed to aid IT in dealing with the complex data protection issues found in a sophisticated VMware® vSphere™ 4 environment supported by host

servers running either ESX® or ESXi™. In such an environment, IT must deal with multiple layers of abstraction with virtual machines (VMs) running on a variety of hosts as it balances the constructs of RTO and RPO. To deal with these issues, Veeam provides IT with a dedicated image-level data protection package that unifies backup and replication in a VMware environment and seamlessly integrates with existing general purpose backup packages.

A key characteristic of a virtual environment is the encapsulation of VM logical disk volumes as single physical disk files. This representation makes image-level backups faster than traditional file-level backups and enhances restoration as virtual disks can be restored as either a whole image or individual files. That's why many general purpose backup packages integrate with VMware Consolidated Backup (VCB) to provide imaging-based backup. Veeam Backup & Replication, however, can directly leverage the new vStorage APIs in a vSphere 4 environment to improve performance without VCB integration. That means small and medium business (SMB) sites without a Storage Area Network (SAN) in place can leverage the ESXi hypervisor with minimal investment. More importantly, Veeam Backup & Replication software extends disk imaging operations to include on-site or off-site replication for rapid recovery.

To provide IT with a better RPO, Veeam Backup & Replication will by default ensure transaction consistency with either VCB or the vStorage API. In addition, a backup administrator can configure Veeam Backup & Replication to use the Windows Volume Shadow Copy Service (VSS) to ensure transaction consistency for VSS-aware applications running within a VM, including Active Directory, MS Exchange, or SQL Server. What's more, Veeam Backup & Replication utilizes the new Changed-Block Tracking feature of VMFS to accelerate incremental backup and replication, recognizes virtual disks with thin provisioning, and can leverage hot-add for virtual disks, when deployed in a VM.


Veeam Backup & Replication also fully supports the VMware ESXi hypervisor, which many sites are now bringing into production as more servers are bundled with ESXi firmware. Nonetheless, the use of ESXi hosts can complicate backup and recovery processes as the architecture of ESXi differs significantly from ESX. To keep the ESXi hypervisor as compact as possible in order to be able to embed this hypervisor in firmware, VMware does not implement a service console in ESXi. Typically data protection applications, such as backup and replication, utilize the service console in ESX.

Veeam Backup & Replication uses VMware Application Programming Interfaces (APIs) to access ESXi remotely and enable backup and restore of VMs running on ESXi servers over the network. With the introduction of Veeam Backup & Replication 4.1, system administrators can also replicate VMs to, as well as from, an ESXi host. As a result, Veeam Backup & Replication simplifies sophisticated IT backup and recovery procedures in virtual environments that include both ESX and ESXi hosts—a capability that is precisely what IT needs to meet Service Level Agreements (SLAs) with aggressive RTO and RPO requirements.

# Data Protection Test Bed

“For testing performance and functionality of backup and replication processes, we set up two groups of server VMs: eight VMs for backup testing and four VMs for replication between ESX and ESXi hosts.”

## VIRTUAL MACHINES, VIRTUAL TAPES, AND SYNTHETIC BACKUP

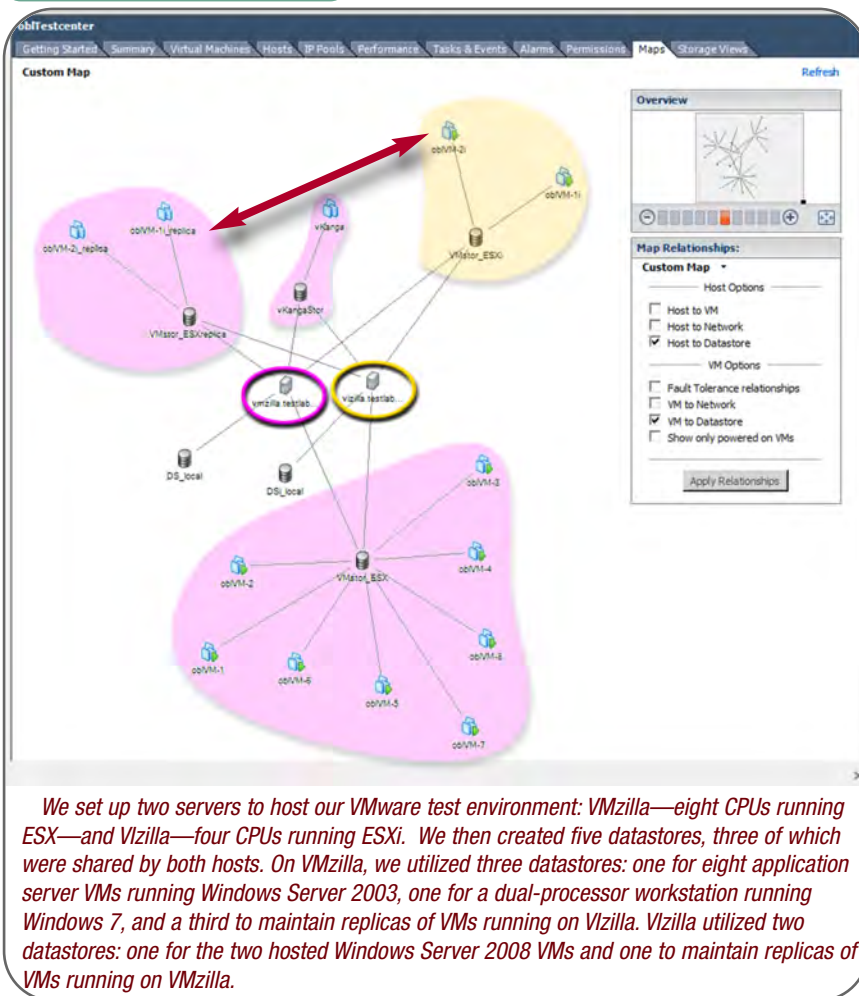


To assess the ability of Veeam Backup & Replication 4.1 to simplify backups and enhance IT operations for disaster recovery, openBench Labs set up four servers to host a data protection scenario based on VMware vSphere 4. Two servers hosted the VMware virtual environment—one ran the ESX hypervisor and the other ran ESXi. All hypervisor datastores were created on a Xiotech Emprise 5000 and each host shared SAN access to every datastore.

We ran Windows Server 2003 on the other two servers and made them members of our Testlab domain. We then installed VMware vSphere 4.0 on one domain server to configure and manage the VMware environment and Veeam Backup & Replication 4.1 on the

other server along with shared SAN access to all of the logical volumes used by the two VMware hosts. We also installed VCB on the server running Veeam Backup & Replication to measure any performance advantages garnered through integration via the vStorage API rather than VCB.

### VSPHERE VM TEST TOPOLOGY



For testing performance and functionality of backup and replication processes, we set up two groups of server VMs: eight VMs running Windows Server 2003 on the ESX host and two dual-processor VMs running Windows Server 2008 on the ESXi host. All VMs in each group were configured as application servers running Windows Server 2003, SQL Server, and IIS. In addition, we set up a VM as a dual-processor workstation running Windows 7 and provisioned it with multiple drives.

For each VM, we utilized thin provisioning on system drives, which were virtualized as vmdk files in a VMware datastore. We stored all work files, however, on Raw Device Mapping (RDM) disks, which were configured in virtual compatibility mode. An RDM volume is a logical disk volume that is formatted by

the VM with a native operating system (OS) file system—NTFS in our tests. That makes an RDM volume directly accessible by other physical, as well as virtual, systems running the same OS. Through the creation of an optional vmdk mapping file by the VMware host, an RDM volume remains manageable through VMFS; however, virtual compatibility does not extend to the implementation of Changed-Block Tracking.

To maximize the execution of end-to-end data transfers over our SAN, we wrote all backup images to a local logical disk that was provisioned from Xiotech Emprise 5000. Veeam Backup and Recovery also provides a virtual appliance mode that works within a VM and avoids network traffic within the VMware environment by leveraging the

hypervisor's hot-add capability with respect to virtual disks. Veeam Backup & Replication uses the vStorage API to attach disks that belong to the VM being backed up to the VM acting as the appliance. In that way, all data is transferred using the hypervisor's I/O stack without having to invoke the network stack.

To lower the amount of disk space required to store backup images, Veeam Backup & Replication 4.1 provides options for data compression, job-level inline data deduplication, and synthetic backup. Within a job policy, a system administrator is able to configure the level of data compression and set the implementation of data deduplication for that job. It's important to note that Veeam Backup & Replication only applies data deduplication within a job: There is no global store of unique blocks. That means two independent backups of the same VM will create identical backup image files.

That deduplication scheme provides a number of interesting performance and functionality tradeoffs. First, the processing overhead for inline data deduplication is dramatically reduced. What's more, a Veeam Backup & Replication image for a job with data deduplication can be transferred without data re-hydration to any off-line media and restored on any server running Veeam Backup & Replication software.

More importantly, the Veeam's data deduplication scheme is transparent when integrating Veeam Backup & Replication with a general purpose enterprise backup package. Integration with an external package is simplified for administrators via a menu option to execute an external command or trigger a batch processing script at the end of the Veeam Backup and Replication process. To test this capability we installed Symantec NetBackup 6.5.4 on our server and a Spectra Logic nTier v80 Deduplication appliance. We used this configuration to assess how an external data deduplication resource could be best integrated with the Veeam Backup & Replication software.

Also dramatically reducing storage requirements and the time needed for a backup window is Veeam's automated synthetic backup technology, which administrators can apply when backing up or replicating virtual machines. Using Veeam's automated synthetic backup, a full backup only needs to run once. After an initial full backup, subsequent backups default to incremental, which can be made more efficient by leveraging the new Changed-Block Tracking mechanism in VMware hypervisors.

Unlike traditional incremental backup schemes, however, Veeam automatically writes the incremental data into the existing backup file (.vbk) to create a new synthetic full backup and then generates a reversed incremental rollback file (.vrb) containing the original data overwritten in the previous full backup. As a result, an administrator can immediately restore the most recent backup. If an earlier backup is needed, the Veeam software will automatically invoke the proper rollback files. What's more, IT administrators can insert a physical full backup into the backup job rotation at any time to comply with any local IT security policy. When this is done, Veeam resets the chain of rollback files to use the new full backup.

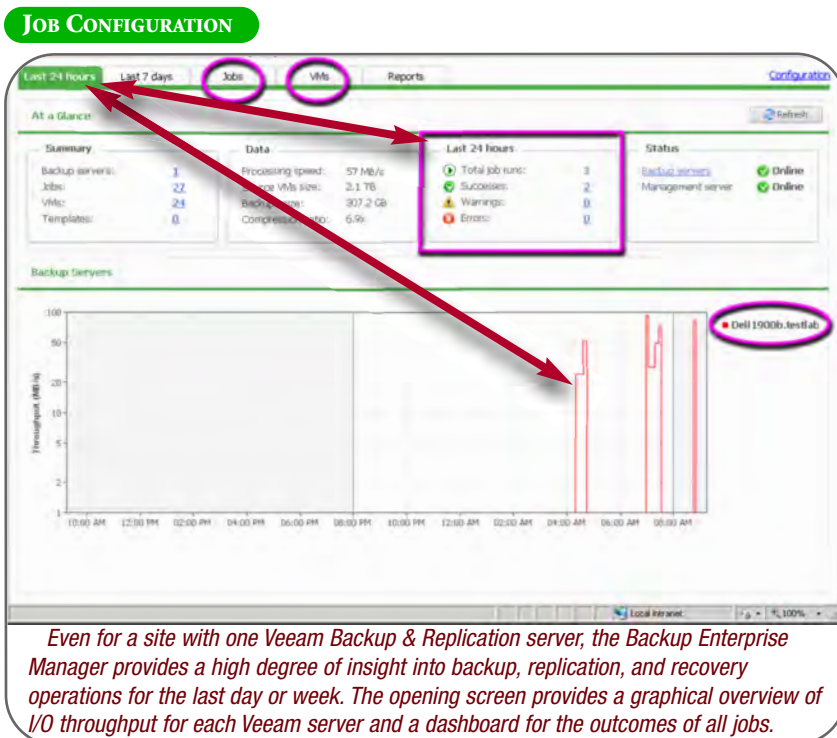


# Backup & Replication Performance

“Not only did we back up eight VMs in sequence more quickly with Veeam Backup & Replication than eight VMs in parallel with NetBackup, the storage footprint of our full backup of eight VMs using optimal compression and data deduplication consumed only 30 percent more space than the backup image of one VM with NetBackup.”

## VM BACKUP AUTOMATION

For CIOs, the rationale for establishing a virtual operating environment rests on improved resource utilization and streamlined IT management. An important outcome is the garnering of the same system reliability, availability, and scalability that typifies large data centers without incurring the operating and capital costs associated with the physical infrastructure of a large server farm.



To help ensure success within a sophisticated VMware environment—especially a distributed environment with multiple backup servers—Veeam Backup & Replication includes two management graphical user interfaces (GUIs). The two GUIs, Veeam Backup Enterprise Manager and Veeam Backup and FastSCP, allow IT departments to fine tune their own balance of centralized and decentralized management and reporting.

Veeam Backup Enterprise Manager acts as a single management point in the distributed enterprise environment by rolling up all of the activities relating to jobs and

VMs across all of the servers running Veeam Backup & Replication. With Veeam Backup Enterprise Manager, IT managers can immediately evaluate what has taken place with respect to data protection processes over the last seven days or the last 24 hours within a VMware environment. What's more, they can clear and resolve issues on any remote server by taking control of backup and replication jobs from a central location.

## JOB CONFIGURATION

The screenshot displays the Veeam Backup and FastSCP GUI. On the left, a tree view shows a hierarchy of servers and containers. A green box highlights the 'vms' container, and another green box highlights the 'vms\testlab\openbench.com' container. Red arrows point from the 'New Backup Job' wizard dialog box to the job table, indicating the configuration process. The job table lists various backup jobs with columns for Name, Type, Status, Progress, and Date. The 'New Backup Job' wizard is open, showing the 'Virtual Machines' section with a table of VMs to be backed up. The table has columns for Name, Type, and Size. The 'Add Objects' dialog box is also visible, showing a tree view of objects to be selected for backup.

*We managed all of the backup and replication tests for our VMware environment using the Veeam Backup and FastSCP GUI. By establishing one connection with our vSphere server, we were able to integrate with all of the ESX and ESXi hosts in our environment. What's more, we were able to create highly flexible job definitions by using VMware container objects, including hosts and datastores to create jobs that would automatically recognize and adjust to changes in VM clients, which is particularly important when vMotion is used at a site for load balancing.*

Complementing the centralized management and reporting paradigm of the Veeam Backup Enterprise Manager, the Veeam Backup and FastSCP software provides a local wizard-based management and reporting application for decentralized operations. At the foundation of the Veeam Backup and FastSCP interface are four key constructs: job definitions, sessions that are instances of job definitions, backup image files that are created by job sessions, and replicas of VMs that are created by replication jobs.

In addition, there are five important wizards: Backup, Replication, Restore, VM Copy, and File Copy. All of these wizards provide simple check-off menus to utilize compression, data deduplication, and synthetic backup for improved processing. Another important aspect of these wizards is the ability to choose VMware container objects to define backup and replication clients. Using these objects allows job definitions to automatically adjust to changes in the number and location of VMs in a VMware environment. Using the Veeam Backup and FastSCP wizards, we began our assessment of Veeam Backup & Replication with the job of backing up the eight server VMs that were running on our ESX host server.

## PHAT ISSUES FOR THIN PROVISIONING

For all of our tests, we used VMware version 7 VMs configured with thin provisioning. To set up performance baselines, we first ran a series of full backups with different levels of compression, settings for data deduplication, and methods of integration with VMware—vStorage API or VCB. These tests provided an eye-opening perspective on the process of backing up a VM.

We began our assessment by running full backup jobs of individual VMs without utilizing either compression or data deduplication. In addition to providing a perfect baseline for our general assessment, this configuration matches the best profile for IT sites planning to stage Veeam Backup & Replication backup images to another repository with a more sophisticated data deduplication scheme, such as that provided by our Spectra Logic nTier v80. What's more, this configuration also provides a good platform

on which to compare Veeam Backup & Replication with an extension of a general enterprise backup application, such as NetBackup to a VMware environment.

We started testing by running backups of single VMs. Our principle metrics were wall clock time and the size of the resulting backup image. Wall clock time replaced throughput as the primary metric because of the manner in which Veeam Backup & Replication handles a backup.

Unlike most competitive backup packages, which utilize VCB to integrate with VMware, Veeam Backup & Replication does not copy the VM files from a datastore on the VMware host to a temporary directory on the Windows server before backing up the files. Veeam reads the VMware data and writes the backup image directly to storage on the Windows server. While NetBackup runs that first step at upwards of 500 MB per second, the second stage slows significantly as the data is reformatted and saved as a NetBackup backup image.

In addition, Veeam Backup and Replication is able to leverage the Changed-Block Tracking that was introduced in vSphere 4. That makes the perceived throughput of an incremental backup appear to be as high as 500MB per second. As a result, wall clock time becomes the best metric for assessing the overall throughput.

For disk-to-disk (D2D) backups, the size of backup images is a critical metric. Standard backup rotation schemes for creating and storing daily, weekly, monthly and annual backups typically increase storage requirements by a factor of 25 times over the original data. The ability to address the issue of storage utilization with compression and deduplication has made these features essential enterprise options.

Thin provisioning of storage is another important way to increase storage utilization. Thin provisioning uses storage virtualization to present a computer OS with what appears to be a fully provisioned volume, while real storage capacity is allocated only when consumed. That makes it important for a backup package to recognize the thin provisioning scheme introduced with vSphere.

<b>Backup Window</b>				
<b>Single VM System Disk Full Backup — No Compression or Deduplication</b>				
<b>Backup Program</b>	<b>Integration</b>	<b>Wall Clock Time</b> hh:mm:ss	<b>Throughput</b>	<b>Image Size</b>
<b>Veeam Backup &amp; Replication</b>	<b>vStorage API</b>	<b>0:02:15</b>	<b>189MB/sec</b>	<b>7.6 GB</b>
<b>Veeam Backup &amp; Replication</b>	<b>VCB</b>	<b>0:03:20</b>	<b>127MB/sec</b>	<b>7.6 GB</b>
<b>NetBackup 6.5.4</b>	<b>VCB</b>	<b>0:13:38</b>	<b>56MB/sec</b>	<b>26.4 GB</b>

On our server VMs, we used thin provisioning for the system disk: Typically each disk consumed 8 GB of the allotted 25 GB. Only the Veeam software recognized this configuration. As a result, Veeam backup images created using either the vStorage API or VCB were

less than one third the size of the NetBackup image and the wall clock time for the NetBackup backup session was six times longer than the Veeam backup session.

<b>Storage Utilization</b>				
<b>Single VM Full Backup — Optimal Compression and Deduplication</b>				
<b>Deduplication</b>	<b>Integration</b>	<b>Wall Clock Time</b> hh:mm:ss	<b>Throughput</b>	<b>Image Size</b>
No	VCB	0:04:19	99 MB/sec	3.5 GB
No	vStorage API	0:03:05	138 MB/sec	3.5 GB
Yes	VCB	0:04:30	95 MB/sec	3.5 GB
Yes	vStorage API	0:03:12	133 MB/sec	3.5 GB

With compression set to its default level of optimal, we tested the effect of data deduplication and compression on a backup of a single VM. In these tests, integration with the vStorage API always provided close to a 27 percent advantage in reduced wall clock time. More importantly, optimal compression consistently reduced the data written to the backup image by 2 to 1.

In our high-speed Fibre Channel environment, there were no throughput bottlenecks with respect to the volume of data being backed up. As a result, the additional CPU processing to provide a 2-to-1 reduction in data via compression was reflected in a modest increase of 18 percent in the wall clock time of a full backup. On the other hand, when throughput was constrained by a gigabit Ethernet connection in a replication process, which requires VM data be transferred to the new host server over a Local Area Network (LAN) so that the new host can create the proper structures, the reduction in the amount of data transferred reduced wall clock time. Reduced wall clock time is also a characteristic of an iSCSI SAN or a LAN based backup, which is made possible by Veeam's support of the vStorage API.

On the other hand, the Veeam inline data deduplication scheme provided no advantage when applied to a single VM with thin provisioning, since we had eliminated multiple blocks of blank data by reducing disk utilization on the host ESX server via thin provisioning. To leverage deduplication, we needed to change our backup strategy from parallel backup jobs of individual VMs to a single job that backed up all of the system disks—approximately 70 GB of data on 200 GB of thin-provisioned storage—of our eight VMs sequentially.

<b>Storage Utilization</b>				
<b>System Disk on Eight VMs Full Backup — vStorage API</b>				
<b>Deduplication</b>	<b>Compression</b>	<b>Wall Clock Time</b> hh:mm:ss	<b>Throughput</b>	<b>Image Size</b>
No	None	0:23:30	145 MB/sec	68.9 GB
No	Optimal	0:30:11	113 MB/sec	38.3 GB
No	Best	0:44:48	76 MB/sec	36.9 GB
Yes	Optimal	0:40:54	83 MB/sec	31.8 GB
Yes	Best	0:51:16	67 MB/sec	30.7 GB

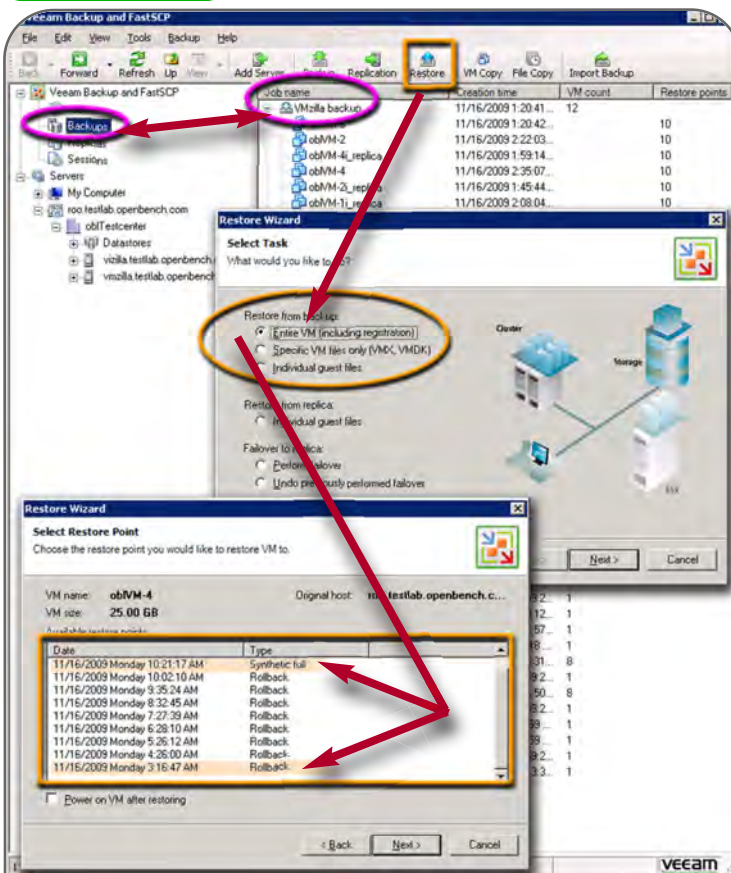
When we backed up the system disks on eight server VMs sequentially in a single job, the deduplication module in Veeam Backup & Replication was able to compare data blocks across all eight VMs to lower the storage footprint of a backup image by roughly 20 percent more than compression alone. Not only did we back up eight VMs in sequence more quickly with Veeam Backup & Replication than eight VMs in parallel with NetBackup, the storage footprint of our full backup of eight VMs using optimal compression

and data deduplication consumed only 30 percent more space than the backup image of one VM with NetBackup. Increasing compression to the highest level that Veeam supports in combination with data deduplication further reduced the amount of storage consumed by 5 percent, but increased processing time by 25 percent.

The 31.8 GB of storage required for the full backup image of the system disks of eight VMs was comparable to the results measured when we backed up the initial 71.6 GB set of Veeam images to the Spectra Logic nTier v80 Deduplication appliance with NetBackup—27.1 GB. Nonetheless, deduplication with the nTier v80 would dramatically improve storage utilization over time, as the nTier v80 maintains a global store of unique data blocks for comparisons with future full backups.

### SYNTHETIC REALITY

#### RESTORE WIZARD



*Using the restore wizard, we first chose how the backup should be restored—as a working VM, selected VM container files, or files of the logical VM. We then restored the VM from the ESX host to the ESXi host and automatically registered the new guest VM. We restored the current version in 22 minutes and the oldest version (eight automatic roll backs) in 25 minutes.*

With a full backup in hand, we could now begin to leverage Veeam by never again doing a full backup. For on going backup sessions, Veeam Backup & Replication defaults to a synthetic backup process. After an initial full backup, Veeam Backup & Replication defaults to performing incremental backups of changed data, which is even more efficient when vSphere Changed-Block Tracking is enabled.

What characterizes a synthetic backup process is the automatic roll up of each incremental backup into a new full backup—dubbed a synthetic full backup—so that the most recent backup is always a full backup ready to restore. In a traditional incremental backup scheme, the full backup is the oldest file. IT administrators must first restore the full backup and then sequentially roll up the incremental files to the required point in time. Since the most recent backup is often the required target, the Veeam synthetic backup process automatically rolls up each incremental backup as it finishes.

In particular, Veeam Backup & Replication creates a new synthetic full backup—a .vbk file—and a roll back service—a .vrb file—from the incremental backup to reverse the traditional restoration process. Using the restore wizard, an IT

administrator can restore a backup to a working VM, selected VM container files, or files of the logical VM. We then restored the VM from the ESX host to the ESXi host and automatically registered the new guest VM. We restored the current version in 22 minutes and the oldest version (eight automatic roll backs) in 25 minutes.

administrator picks a point in time for the restoration point and the Veeam software automatically handles any roll backs.

We tested the process by restoring backups of a VM running on our ESX server as fully functional VMs on our ESXi server. We started with the most recent full synthetic backup. It took 22 minutes to automatically restore the VM, register the VM on the ESXi host, and power on the VM. We then repeated the process with a restore point that required eight roll backs. This time the process took 25 minutes.

#### REAL STORAGE SAVINGS VIA SYNTHETIC BACKUP

<b>Storage Utilization</b>			
<b>Eight VMs Full Backup — Compression, Deduplication, vStorage API</b>			
<b>Backup</b>	<b>Wall Clock Time hh:mm:ss</b>	<b>Throughput</b>	<b>Image Size</b>
<b>Full</b>	<b>0:40:54</b>	<b>83 MB/sec</b>	<b>31.8 GB</b>
<b>Incremental</b>	<b>0:10:08</b>	<b>379 MB/sec</b>	<b>1.1 GB</b>
<b>Incremental</b>	<b>0:07:50</b>	<b>490 MB/sec</b>	<b>0.1 GB</b>

Equally impressive were the performance levels for backup throughput and storage utilization with incremental backups following the initial full backup. Driven by the degree of change in the data and the similarity of the VMs in the backup session, 80 percent of the incremental roll back files required less than 2 GB of storage in our tests.

It took just 33 GB of storage to hold one full Veeam backup and two additional rollbacks. That 33 GB backup image represented 207 GB of original data—a reduction of better than 6 to 1. In comparison, the nTier V80 Deduplication appliance used 42 GB of unique storage block data to represent the same three backup points—a reduction ratio of 5 to 1. More importantly, to garner that reduction in storage space via data deduplication, we needed to run three independent full backups, which incurred a 200 percent increase in wall clock time. Over time, however, the Spectra nTier Deduplication appliance with its ability to compare multiple backup images, will further reduce storage requirements by typically 30 to 1. As a result, sites required to perform full backups can utilize Veeam Backup & Replication with an external deduplication appliance to maintain a high storage utilization rate.

The results of these tests also have profound implications when it comes to effectively implementing a Continuous Data Protection scheme for VMs. Using vSphere's Changed-Block Tracking feature, we were typically able to generate an incremental backup of one or our VMs in about 2 minutes. What's more, we were easily able to restore and register a VM on a different VMware host. Taken together and packaged as a service, those capabilities represent the heart of a disaster recovery process.

#### REPLICATION AND SYNTHETIC CDP

Veeam Backup & Replication builds on its backup constructs to set up replication for disaster recovery. The wizard for replication utilizes the same menus as the backup wizard, which simplifies setup for an administrator.

In a simple sequence of steps, the administrator first chooses the VMs that will be included in a replication job definition. Next a datastore is chosen on a target host. In this datastore, Veeam Backup & Replication creates a folder dubbed VeeamBackup. Within the VeeamBackup folder, folders are then created for each replica VM. Now the administrator makes the same choices for data deduplication, compression, VM quiescence for transaction integrity, and job scheduling that would be made for a normal backup.

Once the administrator finishes with the wizard, there is one very significant difference in a replication process compared to a backup process. Like a backup process, the Veeam proxy server triggers a snapshot of the files for the VM being replicated and reads those files without engaging the VM or its host server. Nonetheless, the Veeam proxy server must engage the new host server for the replica VM. The new host must create duplicate setting files for the original VM and place these files in the new replica directory along with an initial full backup—dubbed replica.vbk. In addition, incremental backup files will be added based on the backup schedule in the job definition.

Engaging the new host requires that the original data is transferred over a LAN to the new hypervisor. This means replication throughput will be constrained at most sites by a 1Gb Ethernet connection. As a result, the amount of data sent during replication has a major effect on the time it takes to complete the replication operation. More importantly, the time necessary for replication to complete also determines the minimum amount of time we would have to wait between launching replication operations and that directly determines the highest RPO that can be met.

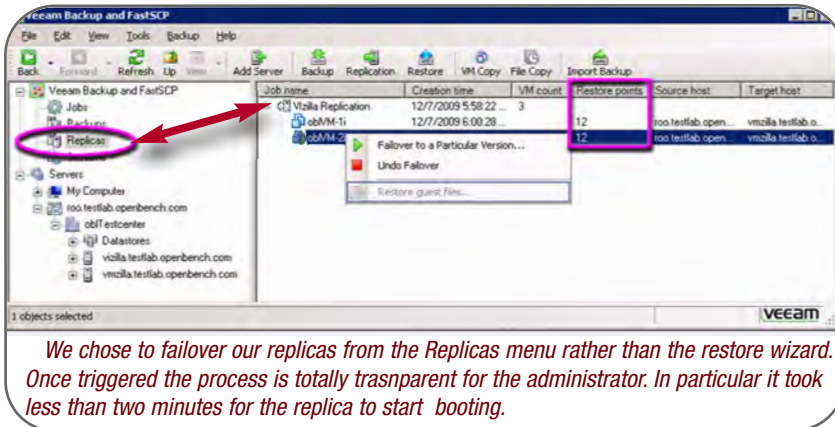
Since replication, even with multiple VMs, is done on a per VM basis, clearly deduplication, just as in our single VM backup test, would not be effective. On the other hand, using optimal compression easily cut network traffic in half. In all of our tests, optimal compression cut data traffic and significantly reduced wall clock time for both the initial full backup and the subsequent incremental backups, which are key to meeting a high RPO.

<b>Replication RPO</b>			
Single VM Replication — ESXi host to ESX host 25GB system volume and 25GB data volume (37GB total data)			
Compression	Backup	Wall Clock Time hh:mm:ss	Throughput
None	Full	0:08:58	67 MB/sec
None	Incremental	0:05:06	116 MB/sec
Optimal	Full	0:07:04	86 MB/sec
Optimal	Incremental	0:03:34	167 MB/sec

In our testing we replicated two VMs running Windows Server 2008 on our ESXi host to our ESX host and two VMs running Windows Server 2003 on our ESX host to our ESXi host. What's more, given the minimal amount of time—about 3 minutes—it took to do an incremental backup with minimal changes, we set up a backup schedule that called for replication every 5 minutes and limited the number of saved replicas to 12 for ESX and 7 for ESXi, which is the

limit imposed by that hypervisor.

### REPLICATION FAILOVER



In effect, we had created a virtual continuous data protection (CDP) process. After allowing sufficient time to populate a number of restore points, we tested the process by corrupting the directory of a VM running on our ESXi host. We then triggered a failover to its VM replica on the ESX server. Within 2 minutes, the replica was booting and within 10 minutes our VM was back online.

At this point the standard option of copying guest files from the replica VM to the original VM was of little use. To rebuild the VM on the ESXi host, we created and ran a standard backup job of the replica running on the ESX host. This created a clean backup without the replica's .vbk and .vrb files. We then restored the backup to the ESXi server and were ready to revert to our normal configuration.



# Data Protection Process Value

“With its ease-of-use features and technical capabilities, Veeam Backup & Replication is perhaps the most advanced solution for VMware backup and CDP currently on the market.”

## PROCESS PERSPECTIVE

Corporate executives think in terms of business processes and expect the services that support those processes to address issues such as availability, performance, security, and business continuity with specific agreed-upon support levels. For IT to create data protection policies and procedures needed to comply with a stringent SLA, IT must establish rigorous data center process control. To help meet these requirements, Veeam Backup & Replication provides two interfaces that support local and central management. This allows IT to tailor their own balance of centralized and decentralized management and reporting.

### VEEAM BACKUP & REPLICATION 4.1: QUICK ROI

- 1) Data protection processes for VMware VMs are policy driven to reduce operator errors and promote automation and include both backup and replication for automated disaster recovery.
- 2) Software compression, data deduplication, and synthetic backup reduced backup data traffic and storage requirements by a factor of 6 to 1 in tests of eight VMware VMs over three backups.
- 3) Rapid restore from single backup can encompass VMware container files, logical files for the guest VM, or can install, register, and launch a new fully functional VM.
- 4) Veeam supports transaction consistency at both the VM level via VMware snapshots and at the application level via MS Volume Shadow Services (VSS) for VSS-aware applications such as Active Directory, SQL Server, and Exchange.
- 5) Synthetic backup foundation provides for near-CDP support with replication, which can be used with VMs on either ESX or ESXi hosts.

With full support for vSphere and synthetic backup technology, which includes in-line deduplication and compression, Veeam Backup & Replication combines both backup and replication protection in one unified package. With a common technology foundation and wizards that employ the same menus for both data protection processes, IT is able to easily implement a robust data protection plan capable of meeting any corporate SLA featuring aggressive recovery point and recovery time options. Using Veeam Backup & Replication, IT can support near-CDP levels of replication for VMware virtual machines with full recovery time measured in minutes.

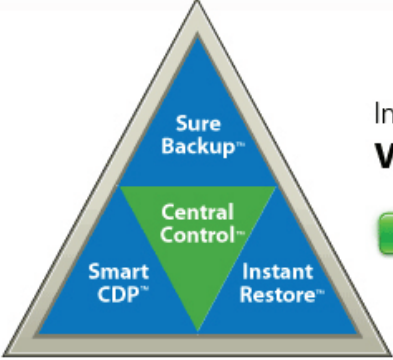
In addition, the Veeam synthetic backup scheme supports continuous incremental backup, which can be modified to include full backups to meet local SLAs. With incremental backups, the data storage footprint of multiple backups can be reduced on the order of 30 to 1. As a result, Veeam Backup & Replication reduce storage costs as it increases the reliability and availability of a VMware environment. With its ease-of-use features and technical capabilities, Veeam Backup & Replication is perhaps the most advanced solution for VMware backup and CDP currently on the market.

# About Veeam Backup & Replication

Veeam Backup & Replication is an evolution of the proven Veeam FastSCP engine, providing unmatched reliability verified by nearly 100,000 VMware professionals over the last three years.


Through continued, rapid innovation in response to customer feedback, Veeam Backup & Replication offers many industry-firsts, including full native support for advanced functionality in VMware vSphere. It helps ensure business continuity in virtual environments and allows customers to reduce costs, minimize risks, and enhance service levels.

## Virtualization-Powered Data Protection™



Integrated with FastSCP  
**Veeam Backup & Replication**

[Learn More](#) [Download](#)



### SureBackup™

- 2-in-1: backup and replication
- 10x faster: vSphere and vStorage
- Proven architecture; FastSCP engine
- SmartDedupe™

### InstantRestore™

- Best RTO and RPO
- VM or file level (Windows and Linux)
- Self-restoring backups
- Advanced VSS support

### SmartCDP™

- Replicate on-site for HA, off-site for DR
- Fraction of cost of traditional CDP
- vSphere Changed-Block Tracking
- Replica rollback

### CentralControl™

- Federation of backup servers
- Reporting and alerting
- Unlimited scalability
- Automation via SDK

**#1 for Virtualization™**  
Data Protection and Management