

WHITE PAPER

# Understanding VMware Consolidated Backup



# Contents

<b>Introduction.....</b>	<b>1</b>
<b>What is VMware Consolidated Backup?.....</b>	<b>1</b>
<b>Detailed Architecture .....</b>	<b>3</b>
<b>VMware Consolidated Backup Operation .....</b>	<b>6</b>
Configuring VMware Consolidated Backup.....	6
Backing up Virtual Machines .....	7
Restoring Virtual Machines.....	7
<b>Troubleshooting .....</b>	<b>8</b>
<b>References.....</b>	<b>8</b>

## Introduction

Virtual machines have become a critical part of many data centers. More and more production applications are run in virtualized environments. The need to protect these virtual machines running the production applications against data loss is urgent. VMware Infrastructure provides many different capabilities to achieve this objective.

Customers use backup methodologies already deployed in physical infrastructure to protect virtual machines. For some environments, this solution is sufficient. However, for the vast majority of customers, installing and maintaining backup software agents inside all virtual machines adds management overhead. This solution is also difficult to scale in consolidated environments in 24x7 datacenters as backup windows shrink and performance requirements become stringent. In general, agent-based backup can be problematic for environments where speed of backup and ease of management of the backup infrastructure are critical.

VMware Consolidated Backup offers a superior data protection solution for virtual machines. Consolidated Backup enhances the ability of backup software to take fast and efficient backups of the virtual machines. It offers a flexible solution by offloading the backup workload from production ESX Server hosts to proxy servers.

This paper explains how the Consolidated Backup technology helps in implementing an effective data protection solution in your datacenter.

## What is VMware Consolidated Backup?

VMware Consolidated Backup is the preferred method of protecting your virtual machines in a VMware infrastructure environment. Consolidated Backup offers an easy way to offload the backup workload to the backup proxy server. The ESX Server components of Consolidated Backup are preinstalled on each ESX Server host, so no separate installation is necessary on the server side. Consolidated Backup software installation is required only on the proxy server.

All major backup vendors fully support using Consolidated Backup to protect virtual machines.


Consolidated Backup offers many benefits for protecting virtual machines in a VMware Infrastructure environment. These include:

- Offloading backup workload from production ESX Server hosts and consolidating it on one or more dedicated backup proxy servers
- Reducing administrative overhead by centralizing backup management on the backup proxies and eliminating the need for backup agents on each virtual machine
- Eliminating the need for a backup window by leveraging VMFS snapshot technology for backup
- Eliminating network traffic on the network by backing up virtual machines over the storage network
- Allowing efficient incremental backups of virtual machines (file-level backup only)
- Allowing online backup of virtual machines (file-level and image backup)
- Integrating with current backup software to back up virtual machines

Consolidated Backup is part of the VMware Infrastructure suite of products. VMware ESX Server is a production-proven virtualization layer run on physical servers that abstracts resources to be provisioned to multiple virtual machines, as shown in Figure 1.

ESX Server uses the VMware Virtual Machine File System for storage of virtual machines. This high-performance file system is optimized to store large files such as virtual machine disk images

(.vmdk files). Each virtual machine contains one or more disk images. When an administrator takes a snapshot of a virtual machine on the ESX Server host, a point-in-time copy of the virtual machine is instantaneously created. Consolidated Backup accesses this snapshot to back up the virtual machine.




**Figure 1: VMware Infrastructure architecture**

Another important element of VMware Infrastructure is VMware VirtualCenter. VirtualCenter manages the ESX Server hosts and virtual machines running on them. VirtualCenter is used to manage and report on the state of all the virtual machines.

As mentioned above, Consolidated Backup can be integrated with all major third-party backup software to provide a nondisruptive backup solution. Backup software in a Consolidated Backup environment continues to handle backup scheduling, reporting, and the data movement between proxy server and secondary storage. Consolidated Backup enables backup software to automate the task of making virtual machine disk images available to the proxy server for backup.

## Detailed Architecture

Consolidated Backup takes the snapshot of the virtual machine disk files and makes it available to the proxy server for backup. Before creating the snapshot, Consolidated Backup flushes the transient writes in the guest operating system and suspends any further writes for a few seconds in order to create a crash-consistent virtual machine image.



**Figure 2: VMware Consolidated Backup architecture**

Once the snapshot is taken, a separate physical machine — the backup proxy — mounts the virtual machine disk image as if it were a locally attached volume so a backup agent running on the proxy can read and back up the volume, as shown in Figure 2.

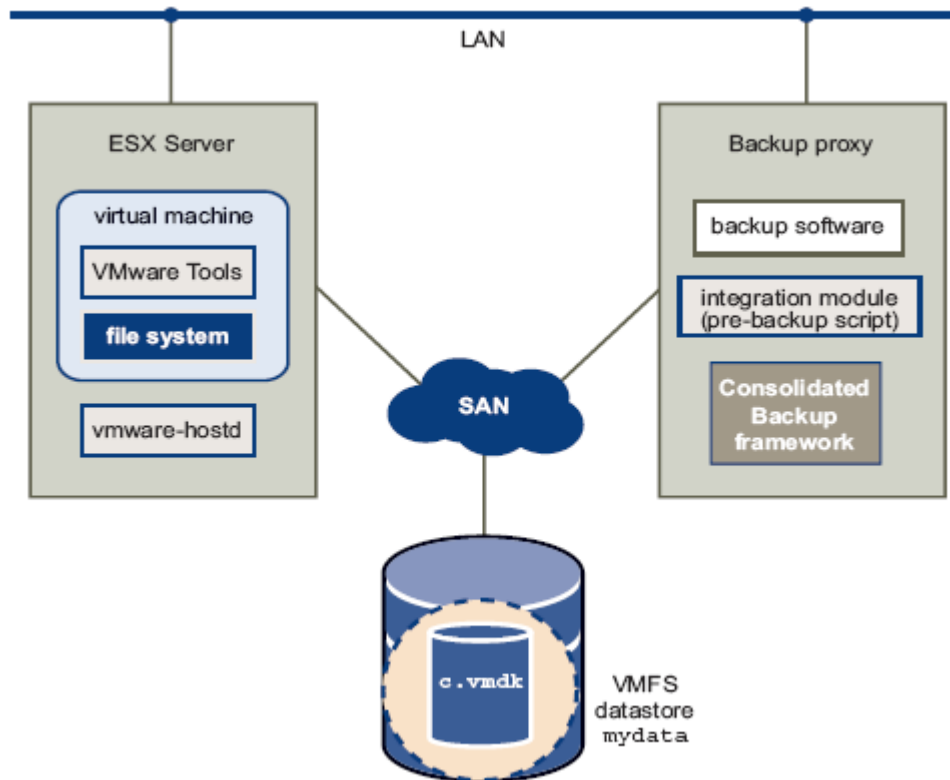
When making an image-level backup of the virtual machine, Consolidated Backup exports the virtual machine disk images to the proxy server. The backup software then backs up these image files normally.

Consolidated Backup also supports backup of virtual machine disks that use Raw Device Mapping or RDM, which allows LUNs to be directly attached to the virtual machine. RDMs are used by applications that need direct access to a raw disk, usually for third-party SAN management software interoperability. Consolidated Backup supports RDMs configured in virtual compatibility mode. For more information on RDMs, see the *VMware SAN Configuration Guide*.

The Consolidated Backup package is installed on a proxy server. This server must be able to access all SAN LUNs on which virtual machines reside and it must be able to communicate with the VirtualCenter Management Server. As mentioned earlier, Consolidated Backup uses the VMFS snapshot functionality of VMware Infrastructure to create a copy of a virtual machine. The biggest advantage of this method is that virtual machines are online and working while the backup is taking place. Also the backup data is accessed on the storage network, not on the production network, so the backup operation is less disruptive.

The Consolidated Backup software has two components: a VLUN driver and, depending on the backup software you are using, an integration module. The VLUN driver enables mounting of virtual machine snapshots on the proxy server, manages snapshots, and makes the mounted drive available for backup.

All major backup software supports Consolidated Backup. Some vendors utilize integration modules to integrate with the Consolidated Backup driver. A few vendors have even tighter integration. Rather than using integration modules, they use an API that provides access to Consolidated Backup functions. For more information on integration modules, see the references section at the end of this paper. Figure 3 shows the components of a Consolidated Backup solution utilizing integration modules.



**Figure 3: Components of VMware Consolidated Backup**

The following steps take place when you back up of a virtual machine with Consolidated Backup:

1. Backup is initiated on the proxy server either by the backup scheduler or manually. The backup agent installed on the proxy server makes a call to the integration module to initiate the Consolidated Backup backup.
2. The Consolidated Backup integration module uses the source virtual machine's information to communicate with VirtualCenter to initiate the backup with the `VMsnap` command.
3. The `VMsnap` command takes a point-in-time snapshot of the virtual machine. All file-system buffers in the guest operating system are flushed to commit writes to the file system. Optionally, application-specific agents can be used to quiesce the applications. For more information on the processing that happens before and after the snapshot is taken, see the *Virtual Machine Backup Guide*.
4. The entire process of taking snapshots lasts only a few seconds. Writes are resumed after the snapshot is created.

All future writes go to a special file, called a delta disk file. This process maintains the state of the

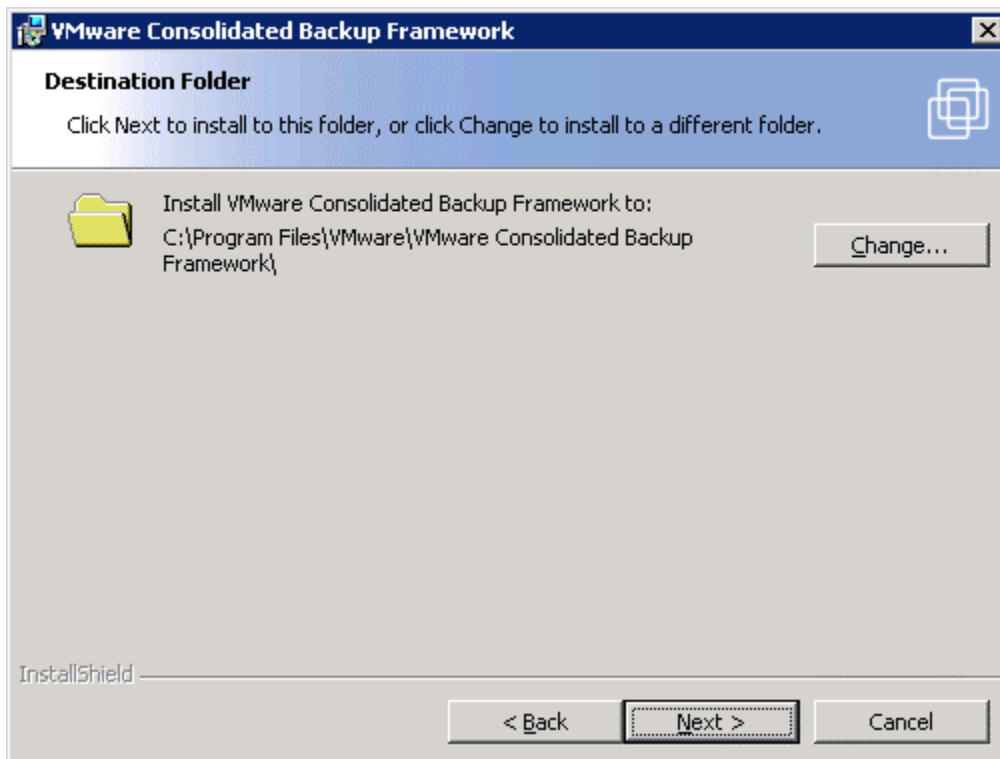
- .vmdk file as a point-in-time copy of the virtual machine. The .vmdk file represents the state the virtual machine was in when snapshot was created.
5. A block list of this .vmdk file is created and sent to the Consolidated Backup proxy server's VLUN driver. A snap ID, which uniquely identifies the snapshot, is created and sent to the proxy server. Backup Information is updated with the snap ID and used by the proxy server for processing the backup.
  6. Using the block list information, the VLUN driver presents a read-only drive to the proxy's Windows operating system. For image-based backups, full virtual machine images are presented as files to the proxy server.
  7. The backup software agent then moves the data from this read-only drive or image file to the secondary storage — for example, a tape library.
  8. After backup data movement is complete, the Consolidated Backup integration module un-mounts the drive and the snapshot is removed by consolidating the delta disk data into the .vmdk file.

# VMware Consolidated Backup Operation

## Configuring VMware Consolidated Backup

VMware Consolidated Backup is a licensable feature of VMware Infrastructure. Consolidated Backup program files are preinstalled on the ESX Server host, and no special installation or configuration is necessary on the ESX Server host itself.

The Consolidated Backup software and integration module are installed on the proxy server. Figure 4 shows the installation on the proxy server. The software is available for download from the VMware Web site or from the VMware Infrastructure installation CD-ROM.



**Figure 4: VMware Consolidated Backup software installation**

The proxy server is typically a physical host connected to the SAN where the virtual machines are stored. The proxy server is also on the same network as VirtualCenter and your main backup server. The backup software agent is installed on the proxy server. For more detailed information on the configuring the proxy server, see the *Virtual Machine Backup Guide*.

After you have configured the proxy server, configure the backup software and its integration module to work with the Consolidated Backup software. More details on configuring the backup software and integration modules can be found in the documentation for your backup software. Some backup vendors provide full UI integrations that eliminate the need for downloading any scripts. This provides the best user experience and is the product direction that VMware encourages. For more information on backup software that works with Consolidated Backup, see the References section at the end of this paper and follow the link for the Consolidated Backup Software and integration modules.



Consolidated Backup is now configured to back up the virtual machines. No special configuration is required on virtual machines to protect them, but you should install the latest version of VMware Tools in the virtual machines.

## Backing up Virtual Machines

You can back up virtual machines with Consolidated Backup more easily by creating a schedule using your backup agent. After you create the schedule, the backup agent works with the Consolidated Backup using the integration module to perform the backup, following the steps outlined in the detailed architecture section.

Backup administrators normally use the backup software interface to initiate or modify the backup. If the virtual machine to be protected is running an application that needs special quiescing steps, Consolidated Backup allows you to run scripts inside the virtual machines before and after backing them up. Using these scripts, you can program application quiescing before the backup and unquiescing after the backup.

Depending on the backup software you use, there are many features available for backing up virtual machines. Refer to your backup software documentation for details on the features it provides. The download page on the VMware Web site lists Consolidated Backup integration modules that are available for download (see the References section at the end of this paper for the URL).

## Restoring Virtual Machines

The easiest way to restore a virtual machine is to use the new VMware product VMware Converter. VMware Converter offers an easy way to restore the virtual machine to the ESX Server host of your choice from a Consolidated Backup image-level backup.

VMware converter reads a Consolidated Backup image-level backup of the virtual machine and recreates a virtual machine at the destination you specify on the ESX Server host. You can set such optional parameters as a new IP address or port group for the restored virtual machine. This makes it extremely easy to recover virtual machines from the backups.

You even have the option of restoring a virtual machine directly from the proxy server to the cluster and making it part of a resource pool of your choice.

In all the above cases, VMware Converter restores the virtual machine and registers it with the VirtualCenter server managing the host or cluster.

If you have created an image-level backup of the virtual machine using a tool other than Consolidated Backup before upgrading to the latest 3.x version of ESX Server — for example, using ESX Server 2.5 — you can restore that image using the VMware Converter, too. For more information on this process, see the VMware Converter user manual.

Even though VMware Converter offers the easiest restore path, other ways of restoring from Consolidated Backup backups are available, as well. For a more thorough discussion of restore options, see the *Virtual Machine Backup Guide*.

## Troubleshooting

This section outlines effective methods for identifying some common errors associated with using VMware Consolidated Backup.

Before you start troubleshooting Consolidated Backup:

- Verify that all prerequisites are met.
- Verify that the configuration meets the storage and networking requirements for all ESX Server hosts, VirtualCenter, and the proxy server.

The VMware Infrastructure *Basic System Administration* guide provides details on setting up networking on the ESX Server host and on setting up VirtualCenter. Setting up shared storage is essential for Consolidated Backup. Make sure the Consolidated Backup proxy server can access LUNs that hold VMware VMFS datastores. For more information on SAN configuration, see the *SAN Configuration Guide*.

A special debugging utility — `vcbSanDbg` — is provided with the Consolidated Backup software. This utility is a Windows program that collects the SAN information as seen from the proxy server. This information is a good starting point to determine whether your storage is properly configured.

When the VMware Infrastructure software reports an error, various logs are updated with diagnostic information. Backup agent logs on the proxy server may also contain information useful in diagnosing errors. The `vcbMounter` output file can contain vital information to help narrow the scope of diagnostic investigations.

You can increase the verbosity level of `vcbMounter` output for diagnostic purposes by adding the `-L 6` parameter to the command line used to back up the virtual machine, as shown below:

```
vcbMounter -h bc.vmware.com -u backup -p XXXXX -a  
\ipaddr:vm1.vmware.com -r d:\vt1\vm1-backup -t fullvm -L 6
```

This command generates output including information messages that are very useful in diagnosing problems.

If snapshot problems are reported in the `vcbMounter` output, one place to look for additional information is the ESX Server host on which the virtual machine is running. On the service console, check `hostd.log` and `vmkernel.log`. These log files are in the `/var/vmware/log` directory.

A special purpose script provided on ESX Server — `/usr/bin/vm-support` — collects and packages all relevant ESX Server system and configuration information and ESX Server log files. This information can be used to analyze the problem you are encountering and to report the bug to VMware support engineers.

## References

For more in-depth technical information on VMware Infrastructure, see <http://www.vmware.com/vmtn>

For VMware infrastructure documentation, including the latest hardware compatibility list and release notes, see [http://www.vmware.com/support/pubs/vi\\_pubs.html](http://www.vmware.com/support/pubs/vi_pubs.html)

To download the VMware Consolidated Backup package and integration modules, see [http://www.vmware.com/download/vi/drivers\\_tools.html](http://www.vmware.com/download/vi/drivers_tools.html)

Revision: 20070926 Item: WP-027-PRD-01-01



**VMware, Inc. 3401 Hillview Ave. Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 [www.vmware.com](http://www.vmware.com)**

© 2007 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,961,941, 6,961,806, 6,944,699, 7,069,413; 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149,843, 7,155,558, and 7,222,221; patents pending.

VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

