

Compliments of:

veeam

VMware Backup

for
dummies[®]
A Wiley Brand



Understanding
VMware backup basics

Recovering
your data

Integrating VMware
backup and cloud

Veeam Special Edition

Lawrence Miller

About Veeam

Veeam® is the leader in backup, recovery, and data management solutions that deliver Modern Data Protection. We provide a single platform for Cloud, Virtual, Physical, Software as a Service (SaaS) and Kubernetes environments. Our customers are confident their apps and data are protected and always available with the most simple, flexible, reliable, and powerful platform in the industry. Veeam protects over 400,000 customers worldwide, including 82% of the Fortune 500 and 69% of the Global 2,000. Veeam's global ecosystem includes 35,000+ technology partners, resellers and service providers, and alliance partners, and has offices in more than 30 countries.

To learn more, visit www.veeam.com or follow Veeam on LinkedIn [@veeam-software](https://www.linkedin.com/company/veeam-software) and Twitter [@veeam](https://twitter.com/veeam).



VMware Backup

Veeam Special Edition

by Lawrence Miller

**for
dummies**[®]
A Wiley Brand

VMware Backup For Dummies®, Veeam Special Edition

Published by:

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2022 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-119-86484-4 (pbk); ISBN 978-1-119-86485-1 (ebk)

- » Getting the basics about virtual machines
- » Understanding virtual machine backups

Chapter 1

Understanding VM and VM Backup Concepts

Server virtualization is one of the most disruptive innovations of the past twenty years. It is a foundational technology in cloud computing and has been widely deployed by businesses of all sizes in both on-premises data centers and in the public cloud. Today, virtual machines (VMs) host practically any workload from mission-critical enterprise applications to dev/test environments supporting DevOps pipelines.

Like any other server in your data center, backing up and restoring VMs reliably and efficiently is a must for business continuity and disaster recovery. In this chapter, you learn about the basics of VMs and VM backups.

Looking at VM Basics

A virtual machine (VM) is created on a physical host such as a desktop computer or, more typically, a server. The host runs a hypervisor (such as ESXi in VMware vSphere) which allows one or more “guest” operating system and application to run at the same time as VMs on the host. The hypervisor manages the dynamic allocation of the host’s hardware resources — CPU,

memory, network interface card(s) (NIC), and disk(s) — to each VM running on the host (see Figure 1-1).

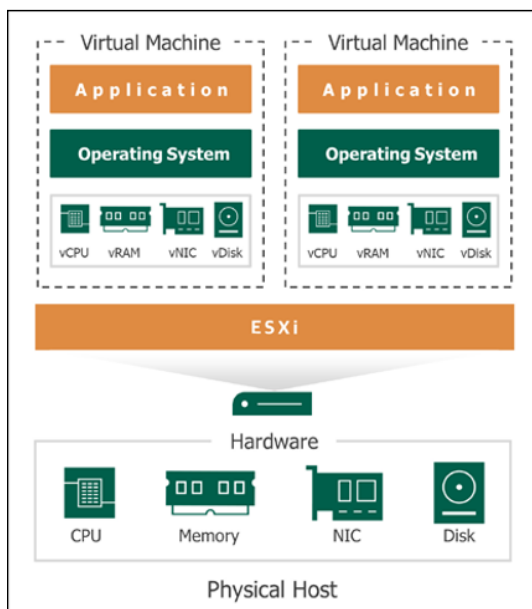


FIGURE 1-1: The hypervisor dynamically allocates hardware resources to VMs running on the physical host.

Some of the most common use cases for VMs include the following:

- » **Workload consolidation:** Multiple application workloads with different requirements can be run on separate VMs to avoid potential interoperability issues between applications running on the same server and to take advantage of unused capacity on a physical server. For example, some applications may only be busy at certain times (such as during business hours or at the end of the month) while others may run primarily during “off” hours (such as nightly batch processing).
- » **Rapid provisioning and dynamic scaling:** VMs can be quickly “spun up” from an image or clone without having to first wait for physical hardware to be purchased, received, and installed in a data center rack. This capability is particularly well-suited for temporary server needs, such as a short-term project or dev/test environment. Additional VMs

can also be automatically “spun up” (scale up) to support periods of peak demand then “spun down” when demand drops (scale down).

- » **Business continuity and disaster recovery:** VMs can be used for business continuity and disaster recovery scenarios to support rapid recovery time objectives (RTOs) without requiring businesses to purchase and maintain physical hardware for a “cold” site.

The business benefits of virtualization include the following:

- » **Increased agility and faster time-to-market** through rapid provisioning and flexible deployment options
- » **Reduced capital investments** for server hardware (particularly when VMs are deployed in the cloud)
- » **Lower operating expenses** for data center space, electricity, and cooling

Introducing VM Backups

VM backups are not the same as physical server backups. For starters, scheduling backup jobs is typically a “set it and forget it” task: You create your backup schedule once, then leave it alone unless there are issues that need to be addressed (such as a job not completing during the backup window) or changes that need to be made. But if a group of VMs all start their backup jobs at the same time and those VMs happen to be on the same physical host during a given backup window (remember, VMs can dynamically move to different physical hosts without human interaction), then there may be trouble afoot! It is possible that you can overwhelm both the hypervisor running on that physical host as well as the hardware resources (CPU, memory, NICs, disk) on that physical host bringing everything to a screeching halt—or at least a painfully slow grind.

Snapshots

You can address some of these challenges with image-level backups, or snapshots. A snapshot captures a point-in-time copy of critical files that comprise the VM. In VMware, these files include the VM configuration files (.vmx and .vmxf) and virtual disk files (.vmdk), among others. However, snapshots by themselves are not backups. There are many key differences between snapshots

and backups, but most importantly, you cannot restore a VM from a snapshot if the original virtual disk is deleted or otherwise unavailable (for example, due to a storage failure).



TIP

Snapshots are typically used to back out failed software updates and other similar use cases. Think of a snapshot as an undo button whereas a backup is more like copy and paste.

Agentless backups

Agentless backups allow VMs to be backed up via a proxy server without installing backup agents on individual VMs. This capability provides many advantages over agent-based backups including:

- » **Less maintenance overhead:** Agents don't need to be installed or updated on individual VMs.
- » **Minimal impact on VMs:** Local resources on the VM are not used for pre-processing and data compression.
- » **Faster backups:** Multiple proxy servers can provide load balancing and multiple VMs (and .vmdk files) can be backed up in a single job.

Backup consistency

Backup (data) consistency is a critical characteristic of a backup solution. There are several levels of backup consistency as follows:

- » **Inconsistent backup ("just okay"):** Essentially, a simple file copy operation that skips system memory and locked files.
- » **Crash-consistent backup ("good"):** Basically, a snapshot of all files on a disk, excluding open files and pending input/output (I/O) operations.
- » **File-level consistent backup ("better"):** All files are backed up with the latest updates, excluding ongoing transactions.
- » **Transaction-consistent backup ("best"):** All data from the same point-in-time is backed up, including completed transactions.



TECHNICAL
STUFF

Data consistency is the sum of validity, accuracy, usability, and integrity of data between applications and across an operating system.

- » Recovering critical workloads
- » Restoring your applications
- » Getting your data back

Chapter 2

Focusing on Recovery

Many may see backup as the most important task in the data center, but remember we back up so that we can recover critical data, workloads, and systems for various reasons ranging from disaster recovery to data reuse that enables business acceleration. In this chapter, you learn how recovery differs for virtual machines (VMs), applications, and data.

Recovering a VM

Recovery is the one thing that requires the flexibility to meet your specific needs with the confidence that you can reliably recover when and where needed. The real question to ask is, “which workloads, applications, and data are most important to my business, and which recovery methods best meet our needs?”

While many recovery methods are available for many varied types of data, applications, and workloads, ultimately you must choose which ones work best to achieve your recovery time objectives (RTOs) and meet your service-level agreements (SLAs).

Recovering a VM from backup requires some planning and testing to ensure you can meet your RTOs and SLAs for a broad array of recovery scenarios. For VMware backup, there are different

levels that can be protected — ultimately based upon your recovery needs. Solutions like Veeam Backup & Replication will provide block-level, image backups that may be leveraged for a full spectrum of recovery options including:

- » **Host-level VMware backup:** Apply VMware backup at the host level to ensure that you capture all of the VMs currently residing on the host as well as all the storage, configuration, and metadata needed for even the most demanding recovery scenarios — including disaster recovery.
- » **Guest-level VMware backup:** Within each vSphere host, you can choose to select individual guests to include or exclude from VMware backup. This method allows you to protect individual guest-level operating systems (OS) and their associated files so that you can perform object-level recovery at the guest OS level.
- » **Changed Block Tracking (CBT):** For all VMware backups, VMware provides a built-in capability that journals all changes made to a guest at the block level. Taking advantage of CBT allows VMware backup products like Veeam to only protect the actual data blocks of what has changed since the last backup, resulting in a much faster, smaller VMware backup.



TIP

Also, take advantage of advanced settings for VMware backup, including VMware Tools quiesce, which will help to freeze the file system during VMware backup. Backup software that supports VMware Tools quiesce for VMware backup will capture all blocks necessary for crash and application-consistent backups.

Comparing Application Recovery and VM Recovery

There are a lot of problems that can arise requiring a system or data to be recovered from backup. Fortunately, there are a lot of recovery scenarios too.



TIP

Be sure you know all of your recovery options and that you're familiar with the process for each option so that you can recover quickly and correctly.

Some problems can be resolved by simply restoring a file. Others may require you to restore an entire VM. Sometimes, just an application item (for example, a SQL database) is all that needs to be restored.

However, there's also a risk of restoring too much. For example, a VM might have multiple SQL databases running on it for various enterprise applications. If only a single application needs to be restored but you instead restore the entire SQL VM, all applications that connect to that SQL VM will be reverted to the restore point along with the single application that needed to be restored. Just restoring the database associated with the single application, instead of all the databases on the VM, will keep everyone else online and uninterrupted.



TIP

Instant VM Recovery is a great capability to have — especially when bad things happen, such as a storage area network (SAN) failure, deletion of a VM, and so on. There's no quicker way to recover a system, especially to a different host or cluster.

Establishing RTOs and SLAs for Data in VM Backups

Many organizations today are considering modernizing their data protection solution. The main catalyst for change: improving RTOs and SLAs, according to the Veeam *2021 Data Protection Report*.

With the confidence that your VM workloads and data are always protected and available, you can achieve new levels of operational excellence and realize business and productivity benefits such as:

- » **Enable digital resiliency.** A modern data protection enables fast and reliable backup and recovery while reducing the overhead and cost of legacy data protection, providing recovery options against ransomware threats, and ensuring organization-wide data protection.
- » **Improve data accessibility and management.** Modern data management increases availability through automated and instant recovery, keeping data platforms healthy without manual intervention, and reducing compliance risks.

Managing data without automated processes reduces efficiency and can increase recovery time. Systems need automated processes to simplify management and recovery while reducing risk.

» **Drive innovation.** Innovation accelerates when data is readily accessible. Legacy backup solutions can trap data in storage silos, not allowing it to be easily re-used. Modern data management can help unleash this data, driving better business outcomes. Backup can become more than just backup. It can be a powerful central repository of data that drives better business analysis, stronger decision making, rapid development, and richer customer experiences.

VEEAM INSTANT VM RECOVERY

VMware vSphere became immensely popular in the late 2000s as a better alternative to deploying costly physical servers in the data center. Unfortunately, backup technology at the time treated VM backups the same as physical backups. This approach quickly ran into serious challenges, most notably, performance bottlenecks. For example, if multiple VM workloads on the same physical host were being backed up at the same time, the physical host's CPU, memory, and network resources often would be quickly overwhelmed causing backup jobs to fail or exceed the time allowed for backup windows.

In October 2010, Veeam changed the world of VM backups forever. Instant VM Recovery in Veeam Backup & Replication can back up an entire VM in less than one minute. This groundbreaking innovation provides a better way to perform a full VM backup (and restore) and is the standard against which all other VM backup and restore solutions are measured today.

Instant VM Recovery allows you to instantly boot a VM directly from the backup repository. The backup repository acts as a Network File System (NFS) datastore that is mounted to an ESXi host. VM data from the repository NFS datastore can be transferred back to the production datastore using either Veeam Quick Migration or VMware Storage vMotion.

IN THIS CHAPTER

- » Taking stock of the complex IT ecosystem
- » Storing your backups in the cloud
- » Using the cloud for disaster recovery
- » Implementing the 3-2-1 rule in the cloud

Chapter 3

Leveraging VMware Backup in a Modern Hybrid Cloud World

In this chapter, you look at complexity in the modern IT ecosystem and its implications for backup and recovery, and how to leverage the cloud for backup storage, disaster recovery, and off-site resiliency.

Recognizing Complexity in the IT Ecosystem

The modern enterprise IT ecosystem is more expansive and powerful than ever. But, with great power comes great . . . complexity! Business infrastructure is deployed around the world in on-premises data centers and clouds of all sorts — including public, private, and edge. Applications no longer run on monolithic three-tier physical servers: Instead, they're installed on virtual machines (VMs) that move dynamically in the data center and in the cloud, or they're built on a microservices architecture that may be composed of thousands of ephemeral containerized services running in hybrid multi-cloud environments.

At the same time, the IT ecosystem has never been more critical. The ransomware attacks against Colonial Pipeline and JBS Foods in early 2021 demonstrated the vulnerability of our supply chains in critical industries and the cataclysmic impact of increasingly prevalent cyberattacks. According to the *Veeam Data Protection Report 2021*, the average cost of downtime is \$84,650. Indirect costs of downtime compound the direct costs and can include:

- » Loss of customer confidence
- » Brand reputation damage
- » Opportunity cost
- » Potential litigation
- » Contract breaches (including service-level agreements)

Business continuity and disaster recovery — underpinned by fast and reliable backup and recovery — have thus taken on renewed importance in the IT ecosystem.

Taking Advantage of Cloud Storage for Backups

As businesses adopt cloud-first strategies, deploying new workloads to the cloud while migrating their existing on-premises workloads to the cloud and re-factoring legacy applications for the cloud, leveraging cloud storage for their backups becomes increasingly attractive.

Simply put, the cloud is the best way to implement an off-site backup strategy. The cloud offers many advantages including:

- » Practically unlimited scalability for growing backup storage needs, including long-term retention for compliance purposes
- » Lower costs for cloud storage services consumed as an operating expense rather than a capital investment
- » High performance and throughput for backups of workloads and data hosted in the cloud, utilizing the cloud provider's network backbone

As part of every organization's ongoing efforts to improve IT data protection, there are five critical stages that follow a progressively more sophisticated data management model:

1. **Backup: reactive recovery** — Organizations at this stage typically use expensive, complex, and unreliable legacy on-premises data protection solutions.
2. **Aggregation: protection across workloads** — Organizations at this stage begin to understand the value of enterprise data and implement cloud-based data protection across multi-cloud environments.
3. **Visibility: from reactive to proactive** — Organizations at this stage have well-protected data and applications and strive to optimize their data use across systems to gain a competitive advantage.
4. **Seamless orchestration** — Organizations at this stage can harness their diversity of data and applications in multiple sites/clouds by automatically moving data/workloads to wherever they generate the most value.
5. **AI-driven automation** — Organizations at this stage leverage artificial intelligence (AI), the Internet of Things (IoT), and the cloud to automatically determine the right time and place to protect, replicate, migrate, and move data and applications.

Making Cloud Backup Part of Your Disaster Recovery Plan

Disaster recovery is often one of the first use cases that businesses undertake when first adopting a cloud strategy. The cloud is a natural fit for disaster recovery because it makes it possible for a business to re-create their on-premises data center on virtual infrastructure at a remote location in the event of a disaster. This approach also has the advantage of not requiring a business to maintain a physical “cold site” data center, hardware, and software in a remote location.

Understanding the Need for Off-Cluster Storage

The 3-2-1 rule of data protection is intended to move your critical backups out of the data center “danger zone” in case primary data services are lost or unavailable. Organizations can rely on

the offsite backup to restore data for operations to another site. Similarly, recording backup data onto multiple media reduces the possibility of media failures or loss; for example, due to a ransomware attack.



REMEMBER

The 3-2-1 rule of data protection states that three copies of data should reside on two different media, with one copy offsite.

Backup to tape, scale-out to object storage, and standby cloud repositories are all viable options for organizations looking to implement the 3-2-1 rule.

VMware COMPONENTS AND INTEGRATIONS

VMware includes many powerful components and integrations to help customers backup their VMs. Here's a quick overview of just a few:

- **Virtual switching** provides a similar concept to traditional networking; however, the switching happens on the VMware solution itself. There are many options with standard virtual switches and distributed switches.
- **vSAN** provides a software-defined storage solution as an alternative to legacy primary storage, making it a great option for mission-critical data storage requirements. vSAN is one of the key building blocks of VMware Cloud Foundation (VCF), which enables vSphere to be run on premises, in the public cloud, and on many private cloud networks.
- **VMware tools** enable important backup capabilities such as determining the IP address or operating system version of a VM and enabling heartbeat and ping tests on a VM.
- **VMware application programming interfaces (APIs)** make it possible to perform application-aware backups for Windows VMs without a direct network connection.
- **vSphere Tags** allow you to include additional information for vSphere inventory objects in order to make them easier to manage, for example, by categorizing objects based on an application type, RTO requirement, or service-level agreement (SLA) metric.

- » Being aware of snapshot challenges
- » Using the right backup agent
- » Sending your backups to the cloud
- » Focusing on your critical applications
- » Keeping your business requirements in mind

Chapter 4

Five Tips for Better Virtual Machine Backups

Here are five tips to help you ensure optimum backup and recovery performance in your VMware vSphere environment.

- » **Don't treat snapshots the same as backups.** Storage snapshots don't replace backups, but they can help minimize data loss in many situations. Some important challenges that need to be addressed in snapshots include:
 - *Storage volume growth* — The storage attached to your VMs can grow very quickly depending on how frequently you create snapshots and how long they are retained.
 - *Performance impact* — As your storage volume grows, the performance of your VMs may be negatively impacted, and if you run out of space on the volume your VMs may come to a crashing halt.
 - *Retention management* — Managing retention policies for different VMs can be cumbersome.
 - *Improper recovery* — If your snapshots aren't application aware, it may not be possible to properly restore a VM from a snapshot.

- » **Use a VM-aware backup agent.** Your backup software should be optimized for virtualized environments to ensure other backup jobs (and production workloads) aren't negatively impacted when backing up a VM, for example, by not backing up every workload on a single physical host at the same time. VM-aware backups are more efficient and will also backup additional information that is needed to restore a VM, such as the VM configuration. Install VMware tools and take advantage of the VMware vSphere application programming interface (API) to enable application-aware image processing and other important backup and restore capabilities.
- » **Leverage the cloud for disaster recovery.** One of the great benefits of a virtualized environment is workload mobility — being able to dynamically move VMs to different physical hosts in the data center or to a different data center altogether. Why not extend that mobility to the cloud? The cloud is an ideal location for off-site storage of backups and has increasingly become a key component of disaster recovery plans.
- » **Ensure application consistency in your backups.** Ultimately, the goal of backing up your VMs is to be able to restore critical applications and services — as well as their associated data — if necessary. Be sure your VMs are in a steady state before backing up your applications to ensure they can be reliably recovered.
- » **Align your backup and recovery strategy to your business needs.** Be sure you understand your business requirements and how recovery time objectives (RTOs) and recovery period objectives (RPOs) for different applications affect your backup schedules and retention periods.

VEEAM

Veeam Platform

Modernize your data protection and eliminate downtime with Veeam® Backup & Replication™, the simple, flexible, reliable and powerful solution for protecting your cloud, virtual and physical workloads.

- ✓ Reduce costs and complexity
- ✓ Any platform. Any storage. Any cloud
- ✓ Ensure bulletproof ransomware protection



Cloud



SaaS



Apps



Virtual



Physical



Data Reuse



Disaster Recovery Orchestration



Monitoring and Analytics



Backup and Recovery



Storage, Analytic and Management APIs



On Premises



Public Cloud



BaaS & DRaaS

Get started for FREE [here](#)

Embark with confidence into VMware backup

Your virtual machines are at the heart of all that your business does, hosting practically any workload from mission-critical enterprise applications to dev/test environments supporting DevOps pipelines. In *VMware Backup For Dummies*, find out how you can implement agentless backups to ensure seamless recovery of critical workloads. Dive in to discover the best recovery methods and data protection solutions to ensure maximum business continuity.

Inside...

- Explore VM and VM backup tools and components
- Implement VMware backup
- Understand the levels of backup consistency
- Use the cloud for backup
- Utilize off-cluster storage
- Discover the top five best practice tips for VMware backup

VEEAM

Go to **Dummies.com**[™]
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-119-86484-4
Not For Resale

for
dummies[®]
A Wiley Brand



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.