

Five Fundamentals for Modern Data Center Availability

David Davis

VMware vExpert

Kirsten Stoner

Product Strategy Coordinator

AVAILABILITY™

for the Modern Data Center

Contents

Introduction	3
Use an availability solution built for modern data centers	3
Select an agentless availability solution	5
Leverage a layered approach to availability	6
Reduce backup data with deduplication	7
Define availability requirements	8
Summary	9
About the Author	10
About Veeam Software	10

Introduction

Company executives, including CIOs and CFOs, have zero tolerance for downtime and data loss. These companies have established high availability requirements for the applications and critical data the organization uses on a daily basis. Sadly, most companies have not found a way to match these expectations with the harsh realities of maintaining the demands of the Always-On Business™. In fact, 82% of CIOs admit to not being able to meet the demand for 24/7 availability of IT services.

Protecting your company's applications and data is more complex than it was in the past and more important than ever. With the introduction of new technologies including virtualization, cloud computing and the advancement of more applications, the aging legacy backup solutions in use today are simply inadequate. This can result in unnecessary risk and pain for companies and IT administrators. This white paper will describe how data protection has changed and how availability solutions are now available that are much more reliable and efficient. Don't take data protection for granted. Continue reading to learn about the five fundamentals of modern data center availability.

Use an availability solution built for modern data centers

There are hundreds of data protection tools out there, but few of them are "virtualization-savvy." Legacy data protection tools tend to see every "server" the same way — as a physical server. By incorrectly assuming that all servers are the same, tremendous inefficiencies occur when you attempt to back up or recover applications and data. For example, lengthy file-based backups are performed when only small blocks of data have changed.

Availability solutions that are built for virtualization can talk directly to the virtual infrastructure. Through this communication, these solutions gain:

- Knowledge of virtual machines (VMs) and the hosts they run on
- Knowledge of virtual storage to see what needs to be backed up
- Knowledge of applications running inside the VM
- Ability to snapshot VMs and back them up with no downtime
- Ability to backup only those VM disk blocks that have changed to drastically reduce backup times and the amount of backup data (referred to as "changed block tracking")

This interaction also becomes a gateway to additional features that you might not expect from your virtual environment, much less from your backup software. For example, you can create virtual lab environments where backups can be automatically tested or used to selectively recover application data. Another example is the ability to recover failed virtual machines from backup, almost instantly, and get end users access to their critical applications. Rarely do legacy backup tools offer innovative features such as these with such immediate business impact.

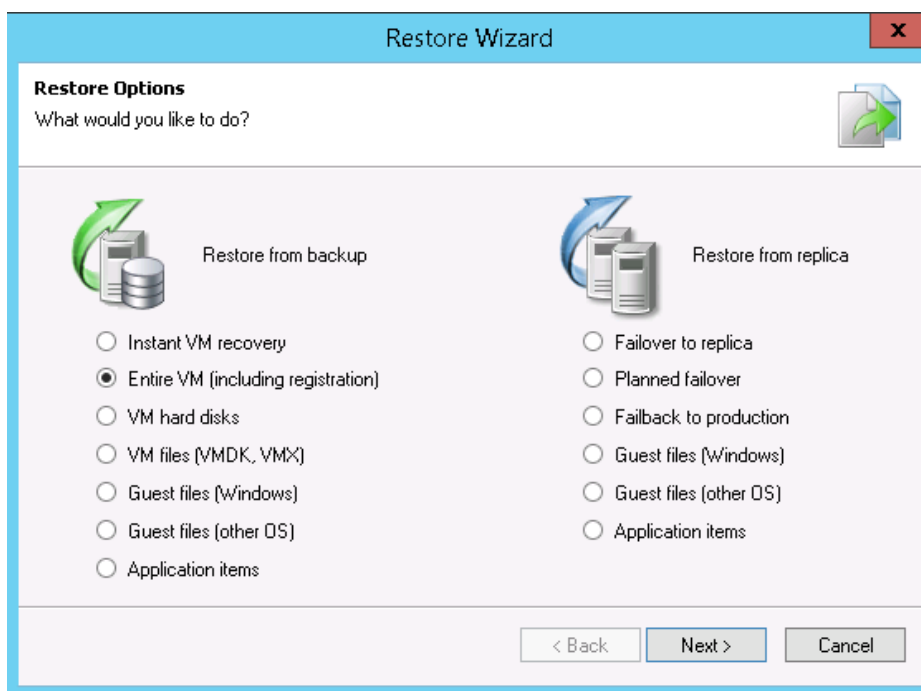


Figure 1: Numerous restore options are available for VMs, including the ability to restore from a backup or a replica. Depending on the software you can also establish failover and failback plans with the use of replication to avoid data loss and downtime.

Over time, some legacy data protection tools have been adapted to recognize virtual infrastructures, but none were originally built with virtualization in mind. With more and servers being virtualized, the smart play is to select an availability solution that is built for virtualization, enabling high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility in your modern data center.

Select an agentless availability solution

Legacy data protection tools require you to install one or more agents on every VM that you need to protect. However, there are downsides to agents, including:

- They require a new piece of software installed on every VM — if a VM doesn't have the agents, it is unprotected and vulnerable to data loss
- They can create conflicts with other applications
- They are difficult to manage, and it's difficult to see which VMs do or do not have agents installed
- They utilize CPU and memory on every VM

The bottom line is that agents are inefficient, across the board.

Some data protection software vendors will say that they are “agentless” because they can do an agentless backup. However, many of these vendors require agents for file-level restore, proper application backup or application data restore. My advice is to make sure that your data protection tool is able to address all backup and recovery scenarios without the need for an agent.

Availability solutions that are built for virtualization can go directly to the hypervisor host or virtual infrastructure management system (vCenter Server or SCVMM), find out the names and locations of virtual disks and then backup those VMs — all without agents.

This means you won't have to install and maintain agents on each VM being protected, your VMs will run more efficiently and you'll gain the maximum benefit possible from your virtual infrastructure.

When it's time to restore a VM or files inside a VM, again, no agents are needed. As you see in Figure 1, there are numerous types of restores available, and none requires an agent.

Leverage a layered approach to availability

Gone are the days of simply backing up your data to tape (and then hopefully storing that tape off site). Modern data centers use a layered approach to ensure availability. This layered approach could describe the 3-2-1 Rule, three (3) copies of data — two (2) copies on different media and one (1) of those copies should be stored off site. Today’s technology offers many different scenarios to establish this approach, including:

- Local backup to disk
- Storage-based snapshots
- VM replication (either on site or off site)
- Archive to tape or cloud storage

The goal is to ensure the availability of your applications and data in as many ways as possible and make restores as fast and easy as possible.






		Strength	Weakness
	Onsite backup 	Streamlined recovery (access backups from disk)	Does not protect against site outages
	Storage-based snapshot 	frequent restore points	Does not protect against storage failures
	VM replica (onsite or offsite) 	Fast recovery (failover to standby VM)	Cost of infrastructure to host
	Offsite backup 	Protects against site outages, long retention	Slower recovery (takes time to retrieve)

Figure 2: A layered approach to data protection leverages new and evolving technologies to achieve the best RTOs and RPOs for all recovery scenarios.

Benefits of a layered approach to availability include:

- Immediate access to backups for instant restore of entire VMs, individual files or application data
- Proven recovery by mounting local backups any time for automated backup verification and disaster recovery testing
- Frequent restore points and ability to meet the most aggressive recovery point objectives (RPOs)
- Automated offsite backup and VM replication for disaster recovery (DR)
- Long-term archiving to tape or the cloud for a final layer of availability for peace of mind and audit requirements

More backup availability?

Follow the 3-2-1 Rule: <http://www.veeam.com/blog/how-to-follow-the-3-2-1-backup-rule-with-veeam-backup-replication.html>

Reduce backup data with deduplication

It's a fact that the cost of data protection will vary greatly depending on the number and size of VMs you have to protect. To optimize your investment in data protection, you must take advantage of technology that reduces the size of your backups. Traditional data protection simply backed up raw data or, if you were lucky, compressed the data.

One of the most common ways to reduce backup size is to use deduplication. Availability solutions automatically perform deduplication as well as compression. Deduplication identifies identical data blocks in source VMs and stores each unique block only once. Because image-based backups used in virtualization capture the entire VM, including the guest operating system (OS), and because the OS is often the same between VMs, there tends to be a lot of duplication. By using deduplication, you can tremendously reduce the size of the backup repository, the time to back up VMs, the amount of backup data replicated off site and the amount of data sent to tape or cloud storage.

Don't attempt to protect your virtual infrastructure without using a tool that includes deduplication. Other space-saving features to be on the lookout for include:

- **Forever incremental backup:** This feature performs an initial full backup but then performs incremental backups, forever. As incremental backups are performed, the full backup is updated with the changes to create a complete backup image of the VM that is always ready to be restored.
- **Support for hypervisor thin-provisioning:** Your backup tool should understand that your hypervisor can create thin-provisioned virtual disks. It should support thin-provisioned disks and maintain thin-provisioning throughout backup, restore and replication operations.
- **Exclusion of unneeded data:** Availability solutions should recognize that your VMs are made up of special files such as the configuration file, swap files, snapshot files and the virtual disk. Not all of those files need to be backed up. You need to be able to specify which of these files you want to exclude from the backup to save time, network bandwidth, and space in the backup repository.

Define availability requirements

The modern data center is all about running 24/7 operations, avoiding any sort of downtime and data loss, and supporting growing amounts of data. It is important to define the correct availability needs for your organization. Data loss costs businesses \$2 million in just one year, and downtime can cost enterprises over \$10 million annually. To be an Always-On Business™, recovery time and point objectives (RTPO™) should be less than 15 minutes.

The modern data center is often defined by being highly virtualized, consisting of modern storage and having a cloud strategy in place. These aspects of an environment can allow for availability and for data centers to have the least amount of downtime possible. Companies may be taking a cloud-based approach to improve availability, whether it is sending data to the cloud or replicating to the cloud. They may be establishing failover and failback plans. They can also be leveraging techniques like storage snapshots or item-level recovery scenarios to avoid the data loss associated with downtime. Whatever approach you choose to take when establishing availability for your data center, it is important to realize that with downtime comes the potential cost of losing money.

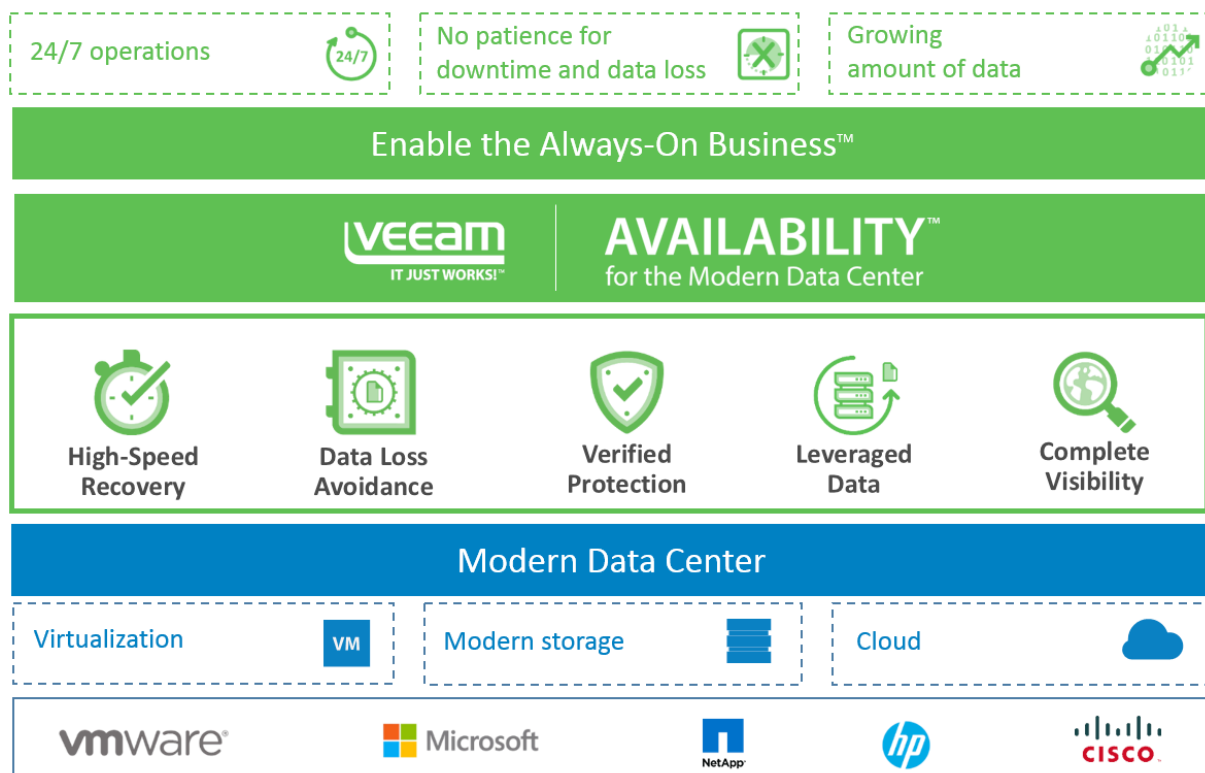


Figure 3. Enabling the Always-On Business with Veeam

Summary

Managing a data center is a complicated endeavor as many metrics, data points and systems are in play. With a constant flow of new applications, new technologies and new problems, it's easy to take core data center responsibilities, such as availability, for granted. In too many instances, data centers stick with "the devil they know" and continue to use the same legacy backup application that they have used for the last 10 years. The problem with that route is that it leaves IT administrators inefficient, end users with more downtime and data loss than necessary and the company with potential availability gaps. Availability solutions know virtualization, are agentless, use a layered approach, include features to reduce the size of backup repositories and support the latest and most popular virtualization platforms. I recommend evaluating availability solutions for use in your data center today!

About the Author



David Davis is a well-known author, speaker, and evangelist on datacenter technologies. David's library of popular virtualization video training courses can be found at Pluralsight.com. He holds several certifications including VCP, VCAP, CCIE, and has been awarded the VMware vExpert award 6 years running. His website, covering the latest in virtualization is VirtualizationSoftware.com and he's a partner at ActualTechMedia.com.



Kirsten Stoner is a Product Strategy Coordinator at Veeam Software. Her responsibilities include community tasks such as providing answers to Spiceworks user's online, social media and content for Veeam campaigns.

You can follow Kirsten on Spiceworks at:

<http://community.spiceworks.com/people/kirsten-veeam>

About Veeam Software

Veeam[®] recognizes the new challenges companies across the globe face in enabling the Always-On Business[™], a business that must operate 24/7/365. To address this, Veeam has pioneered a new market of *Availability for the Modern Data Center*[™] by helping organizations meet recovery time and point objectives (RTPO[™]) of less than 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility. **Veeam Availability Suite**[™], which includes **Veeam Backup & Replication**[™], leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 30,500 ProPartners and more than 145,500 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit <http://www.veeam.com>.

COMING SOON

NEW Veeam® Availability Suite™ v9

RTPO™ <15 minutes for ALL applications and data
Enabling the Always-On Business™
with *Availability for the Modern Data Center™*

To learn more, visit www.veeam.com