Modern Data Protection

# Virtual Backup Strategies: Using Storage Snapshots for Backups


## Hani El-Qasem

Senior System Engineer, Veeam Software

# Contents

# Introduction

Effective data protection is a mandatory element in the modern IT environment. Historically, backup strategies were confined to the last few chapters in an administrator's manual and treated like an afterthought. Now they sit firmly at the forefront of every CIO's mind. The ability to continue business operations after a system failure and the need to fulfil stringent compliance requirements have made backup a necessity—not only for business continuity, but also for business survival. The question organizations need to ask about data protection is not whether to backup their data, but how to backup their data.

IT systems are prone to rapid evolution and present a constantly shifting landscape and the techniques used to protect those systems need to evolve as well. Perhaps one of the most significant changes in recent years has been the advent of virtualization. In the virtual world, legacy backup systems have become unfit for their purpose, causing backup windows to increase beyond a manageable scope. While this paradigm presents new challenges, new opportunities to improve efficiency, cut costs and reduce risks are also created.

This paper will examine the use of storage snapshots as backups for virtual environments. We will evaluate the relative benefits and limitations while also considering where they fit into a holistic backup strategy when compared to a virtual disk-to-disk backup solution such as Veeam® Backup & Replication™.

# Background

Pre-virtualization backup strategies were underpinned by operating system (OS) and application-level features. The typical implementation would involve installing a backup agent into an OS and the agent would be responsible for putting applications into a consistent state for backup; copying backup data across the network to a backup server and subsequently monitoring any ongoing changes.

While this worked well in the physical world, virtualization changed everything as operating systems began to share the same physical hardware. Instead of having one backup agent consuming resources from a physical host, there was an agent for each virtual machine (VM) on that host. This meant that ten agents (based on a 10:1 consolidation ratio) or even more could be contending for the host's CPU, RAM and disk resources. This contention was not only with each other, but also with the applications they were installed to protect. In addition, volumes of data increased to a level where it was no longer feasible to use standard transports to move it across the production network to the backup server. This situation clearly could not continue as virtualization has become the standard practice of datacenters worldwide.

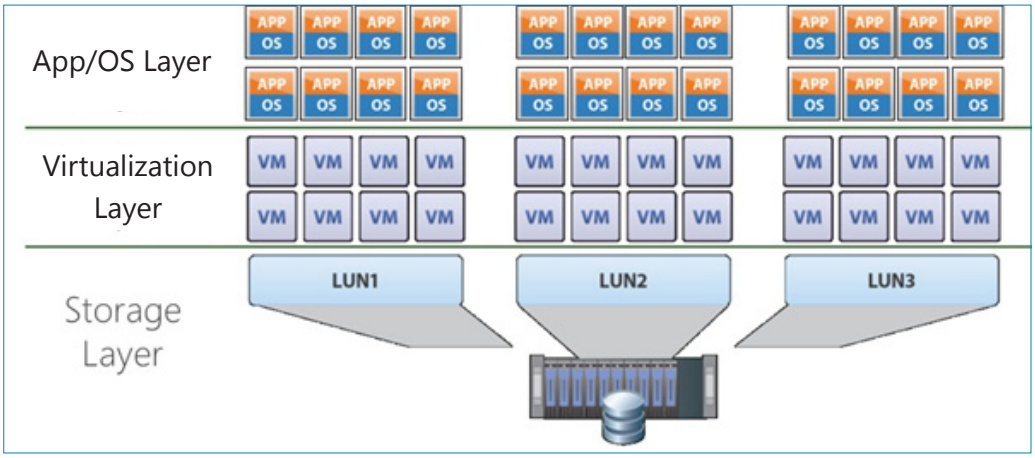


Where virtualization presented new challenges, it also presented new opportunities. The physical world consisted solely of the application/OS layer. The virtual world, while still supporting the application/OS layer, introduced additional underlying layers such as the hypervisor layer (i.e. virtualization layer) and storage layer. These new layers presented many more possibilities for accessing, copying and protecting production data. The hypervisor's facilities allowed the centralized management of these processes without requiring agents inside each VM. The centralized nature of shared storage in the storage layer means there are more effective techniques for moving data from primary to backup storage without having to use the production network.

It is in this context where we'll examine how we might use storage snapshots to protect production data and where a hypervisor-driven disk-to-disk solution such as Veeam Backup & Replication may be more appropriate.

# Storage snapshot overview

In order to understand how we can protect data with storage snapshots, we should first understand what SANs and storage snapshots are. A SAN (Storage Area Network) is a storage device consisting of multiple physical hard disk drives. It would typically be connected to a storage network and managed by software provided by the vendor of the particular model of SAN.

SAN devices are engineered to aggregate disk resources for dealing with large amounts of data. In recent years, additional processing power has been built into the devices in order to offload processing tasks from the physical hosts which are serving up resources to the virtual environment. The basic unit of management for a SAN device is a Logical Unit Number (LUN). A LUN is a unit of storage which may consist of several physical hard disks or a portion of a single disk.

There are several considerations to balance while specifying a LUN configuration. One LUN intended to support VMs running Tier-1 applications may be backed by high-performance SSD disks whereas another LUN may be backed by large, cheap disks and used primarily for test VMs. Once created, LUNs are made available to hypervisors which in turn format them to create volumes (e.g. VMFS – Virtual Machine Files System on VMware, CSV - Cluster Shared Volume on Hyper-V). From this point on, I will use the term LUN and volume interchangeably. A LUN can contain one or more VMs.

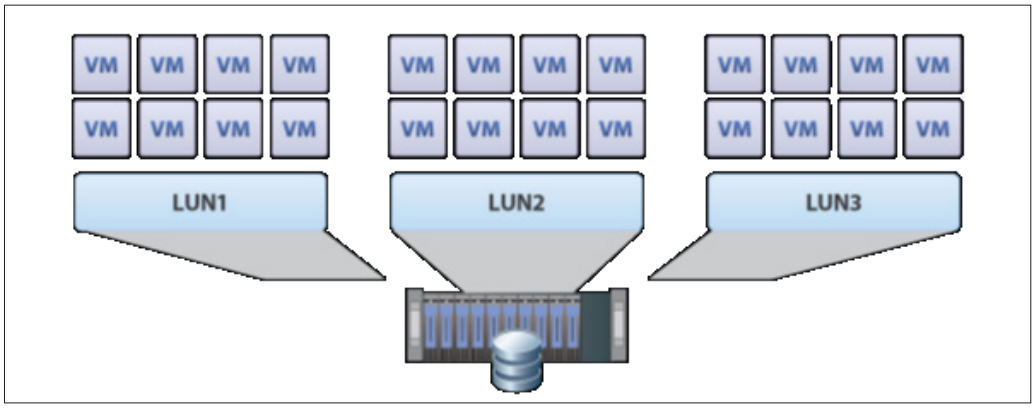


To protect (or backup) a VM, a SAN will create a point-in-time copy of the LUN where the VM resides. The basic mechanism for creating a point-in-time copy of virtual machine disk data is the LUN "snapshot" (a.k.a. SAN or storage snapshot). SANs are able to create LUN level snapshots of the data they are hosting. A LUN snapshot will freeze the entire volume at the point it is taken while read-write operations continue without halting. This should not be confused with a hypervisor-level snapshot (e.g. VMware snapshot) which works higher up the stack.

### Snapshot speed and incremental design

The process of creating a storage snapshot is extremely fast. The mechanism for doing this is specifically engineered at the hardware level, which will in most cases outperform snapshotting using software-based approaches further up the stack (e.g. a hypervisor snapshot). It also removes processing from the host, leaving more capacity for the VMs residing there. This means that it can be done quickly and often presents the possibility of very short Recovery Point Objectives (RPOs). A typical schedule might include taking a snapshot every 30–60 minutes.

At this stage, we should understand what is happening during this process. The first snapshot of a LUN will freeze the entire LUN at the point the snapshot is invoked, creating a base snapshot. The SAN doesn't copy the data to a different location; it simply sits in the same place. Disk writes for that LUN are then redirected to a different area of the disk so those original frozen blocks aren't overwritten.

When subsequent snapshots are taken, the same process of freezing the blocks happens in the new area, but this will only include blocks which have changed since the previous snapshot. We therefore have a process which mimics a traditional incremental backup strategy where a full backup is taken and subsequent backups only take the changes that occur between cycles.

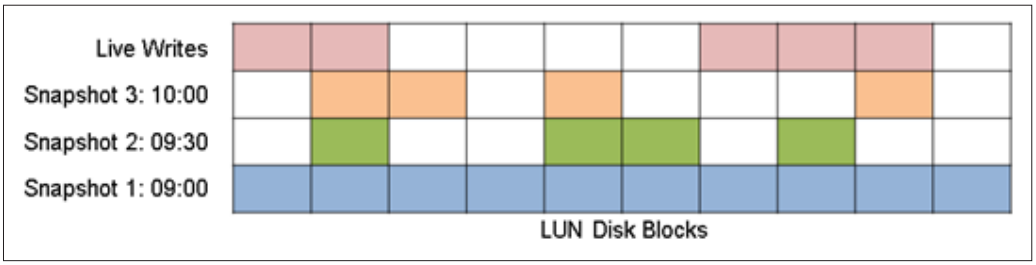


There are obvious benefits to using such a rapid and space efficient strategy

## Recovery from snapshots

During a recovery operation to rebuild an entire LUN at a certain point in time, the SAN pieces together the relevant blocks to make the LUN available as required. For example, rebuilding Snapshot 3 above would look something like this at the block level:

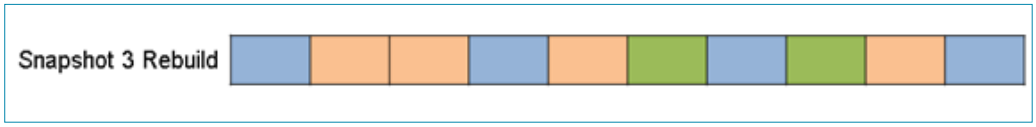


Recovering individual VMs, virtual disks, files or application items is less straight forward. It requires an entire LUN to be recovered and mounted to give access to the more granular data. Beyond this, additional software features are required to interrogate, retrieve and restore these individual items. We'll explore this process in more depth below.

## On-device (primary) vs. off-device (secondary) protection

So far we have discussed what happens to the primary storage system during snapshotting and the subsequent LUN recovery operation. Looking ahead, there is a distinction between the primary on-device protection and the off-device protection.

Data is at significant risk when it's not moved from its original location to a different physical device.To do this, a second SAN from the same vendor (and in most cases the same model) is required before a snapshot replication process is used to move the data "off-device." To perform a recovery in the off-device location, the LUN must be mounted, made available and in some cases copied back to the primary device. This is one of the areas where the process becomes challenging. We'll cover this in more depth below.

Separate storage is the key to ultimate protection for virtualized environments. At Veeam, many of our best case studies come from SAN failures which may include the snapshot logic. In an era of disk-based backups, it's best to provide dedicated storage for data protection separate from VM primary storage.

# Storage snapshot limitations

In theory, we have found an alternative to traditional agent-based backup that is very fast and removes loads from the production environment. Unfortunately, using this approach becomes somewhat more challenging when recovering the storage system alone.

### Linked snapshots on primary storage

First and foremost, we have to ask: "By having the data in the original location via a storage snapshot, is this data fully protected?" We are, in reality, just moving where the disk writes occur with each snapshot and keeping backup data on the environment's primary storage device, which of course will have a massive cost implication. What happens when that backup data starts to fill that device? We'll have to buy more disks and more disk shelves so we may continue to store both primary data and backup data on the same device. While it is possible to use SAN replication to get the data off-device, the process is complex, difficult and adds significant costs to the environment, Especially when considering the volume of data we are attempting to move.

In addition to space implications, the highly interconnected nature of the snapshots (each linked to the previous and ultimately a base image) can also be an area of concern. Minor corruption in any link in the chain can affect all subsequent snapshots, thus significantly increasing risk. Other than basic integrity checking, verifying the recoverability of snapshot-based backups is for the most part non-existent. The linked design also reduces transportability of the backup data.

### Granularity

Using LUNs as a primary unit of management can introduce limitations when more granular control is required at the VM, file or application level. This limitation stems from the storage system's approach to servicing many clients including VMs. Other operating systems may be connected to the SAN.

### Backup consistency

All backup technologies can generally fit into one of the main consistency levels. A consistency level dictates the planned integrity of the recovery point. There are three levels of consistency that can be achieved in any backup process:

- **Crash consistent –** This is analogous to yanking the cord out of the machine supporting the target application
- **File consistent –** The files undergoing file operations are completed and paused before a replication operation is initiated
- **Application consistent –** Application files are put into a consistent state, and anything currently occurring in memory is also committed to disk. Additionally, the applications are made aware that they have been backed up and should continue their operations accordingly.

It makes sense to implement application consistency as this will result in the least possible amount of data loss and the shortest possible restore time in the event of a failure. File system consistency is the second most desirable technique, with crash consistency being used as a last resort.

In order for SAN products to achieve application consistency in a VMware environment, most vendors will use VMware snapshots and the provided VMware Tools' quiescence (i.e. pausing) functions. Unfortunately, whereas VMware Tools puts some applications into a consistent state, it does not provide a truly application-aware implementation. It tells the application to prepare to be copied rather than be backed up. This means that the application is never actually aware of the full backup process and is unable to perform any associated maintenance operations such as trimming transaction logs or placing itself into a state to be restored. Additionally, there are often inconsistencies in the way VMware Tools deals with different guest operating systems making the process even more unreliable.

To negate this problem, SAN device vendors often require VM admins to insert agent software into each VM's guest operating system. While this approach may achieve the required level of consistency, it is complex, difficult, costly and requires careful orchestration of the pausing process among the application, operating system, hypervisor and the SAN. In addition, there are overhead (resources (CPU, RAM), performance and management) in any approach that requires agents to be installed inside VMs, which is why organizations are moving away from agent-based techniques in the first place.

Perhaps the greatest challenge is the lack of granularity with which storage snapshotting occurs. Snapshots are processed at the LUN level. For example,

if there are 30 VMs sitting on the same LUN, you have to put all 30 VMs into a consistent state at the same time. If you can imagine 30 VMs dumping whatever they have in memory to disk at the same time, you may not be surprised that this causes a massive spike in resource utilization and in some cases grinds production hardware to a standstill.

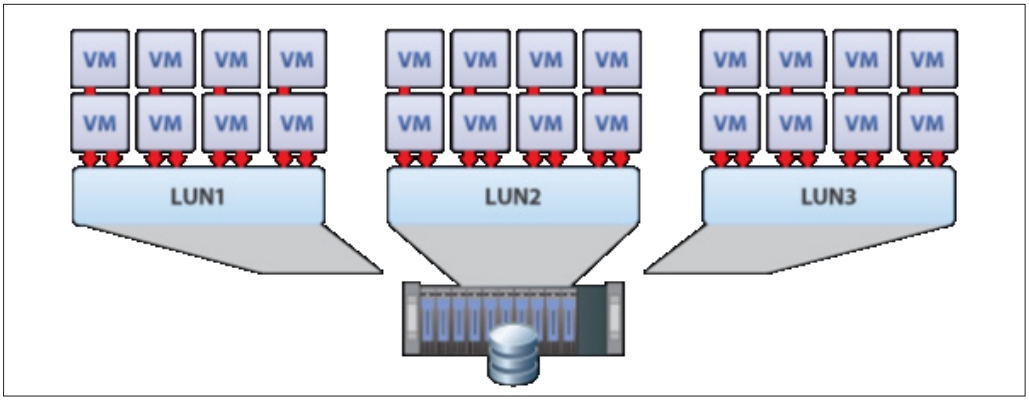


To avoid such issues and spread out the load, separate the VMs onto different LUNs based on the level of consistency they require. What if you already placed the VMs on specific LUNs for performance reasons? Should you subdivide LUNs into multiple LUNs to ensure that both application consistency and performance requirements are achieved? Increasing complexity can only lead to further challenges and problems. It seems like you are being forced to change your virtual environment in order to meet the limitations of the snapshotting process—not a good place to be in.

It should be acknowledged that some vendors state they are able to run the snapshotting process for "single VMs." While this may be true in terms of processing application consistency, the process will still snapshot the entire LUN at the bits and bytes level (i.e. the other VMs will still be snapshotted, but in a crash-consistent state).

### Disk space considerations

SANs are typically designed to manage large quantities of on-device data. Unfortunately, because of the propensity to work with large blocks of data, changes between snapshot cycles are often tracked at a very large block size, sometimes as large as 15 MB. This means that if a 10 KB change occurs on one block and a 20 KB change occurs on another block on the same LUN, a chunky 30 MB of snapshot storage will be required for those relatively small changes.

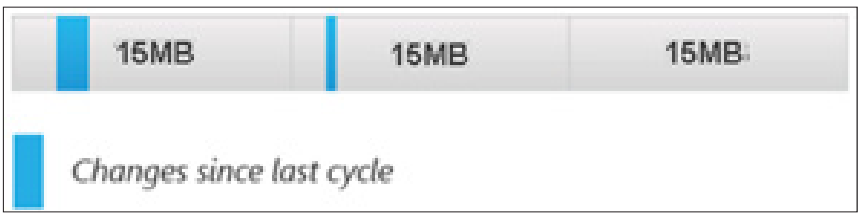

Some vendors have introduced smaller block sizes and deduplication on their primary storage systems to decrease space requirements. However, this doesn't provide the necessary efficiencies to resolve the space issue in most cases. This can create additional issues where heavy I/O is required for the production VMs. All SAN vendors recognize this as a serious limitation of the snapshotting process.

To alleviate the problem, organizations identify data that doesn't need to be backed up and then move this on to LUNs which aren't snapshotted. Examples of data which do not need to be backed up may include in-guest page files, swap files, application log files or other temporary files. While this may sound straightforward, it amounts to a massive redesign and reconfiguration of the entire virtual environment. To separate the relevant data, we need to create additional volumes inside each VM, create additional virtual disks for each VM and then host these disks on different LUNs on the SAN.

When choosing a LUN to host a VM, we need to consider not only performance requirements and application consistency requirements, but also how to split that VM across multiple LUNs in order to exclude certain data from being snapshotted. Not only does this introduce complex management overhead, it also breaks one of the core tenets of virtualization which is "encapsulation." The objective of encapsulation is to bring workloads into VMs. By design, VMs are a set of grouped files which are held together and therefore very portable. Spreading VM data across several LUNs breaks that encapsulation and also introduces many dependencies on the configuration of the underlying hardware, another issue that virtualization was designed to avoid known as hardware dependence.

### Granular restores

When we start to look at more granular restores, the process becomes more complex. TThe following items should be considered least to most granular respectively: LUN snapshot, VMs, virtual disks, files and application items. To restore the most granular item, you must also restore some of the less granular items. An example may look like this:

1. LUN snapshot level: Rebuild a LUN snapshot for a relevant restore point.
2. VM/virtual disk level: Mount the LUN snapshot to make Virtual Machine Disks (VMDKs) available. Use vSphere tools to register and boot a VM or to attach a VMDK.
3. File level: Attach to, or interrogate, virtual disks to retrieve files.
4. Application-item level: Use an application-specific browser to grab an item from the application file. Restoration to the original location requires agents inside the live VM hosting the application.

If the target application item is in an off-device replica, this adds more overhead.

The four steps above simplify the explanation. The process is much more involved and requires additional manual intervention (in some cases running command-line scripts). In recent years, some vendors have been able to provide some automation for levels one and two described above, but support for levels three and four is sadly lacking. Where it is available, it requires legacy agent-based interaction and some installs reportedly take up to four weeks to implement basic capability.

This level of complexity not only increases restore times but also denies organizations the opportunity to delegate restores to less technical employees. This is a serious deficiency that needs to be resolved.

Due to the data being on-device with primary storage, data transfer times can be significantly shortened by using a disk-to-disk backup solution, but this speed may be offset by the complex stages of restoration. So when a low RPO may be achieved during snapshotting, low Recovery Time Objectives (RTOs) may not be easy to accomplish.

### Hardware dependence, moving backups off-device and long-term retention

Common sense will tell us that keeping backups on the same physical device as the data being backed up introduces significant risks. A mechanism must be found to move the data to a secondary device preferably in a different physical location (at the very least a different room or floor in the building).

The only mechanism available for moving snapshot backups is snapshot replication. As stated above, SANs are perfectly capable of processing large quantities of on-device data, however, moving that data off-device is another matter. The fact that SAN technologies are typically not optimized for network transmission and that the smallest unit of management is a LUN snapshot means that it isn't easy to create an effective process to move the data off-device.

The same block size constraints detailed above also have a serious impact on moving snapshots off the primary device, and if a forced restructuring of the virtual environment hasn't already occurred, it certainly will when we try to move snapshots across the network.

Long-term retention is also affected by the heavy hardware requirements to store snapshots. If you are storing backups for multiple years, those secondary SANs are going to require more and more disk. The traditional solution for long-term retention would be to use tape but now the cloud has become a viable target for storing backups. In the case of tape, the bulk nature of snapshots would quickly consume space, and costs would spiral upward. As

for replicating to the cloud, there would have to be another SAN capable of receiving snapshots again. In most cases, this means tying organizations to a vendor in addition to the associated disk space requirements.

# Overcoming snapshot limitations with Veeam Backup & Replication

There are several areas where storage snapshotting provides benefits, but also many limitations which would challenge it's viability as a primary virtual backup solution. Veeam Backup & Replication can address these limitations.

### Linked snapshots on primary storage issue

First and foremost we should state that Veeam Backup & Replication is specifically designed for moving data off primary storage devices in a resource- and speed-efficient manner. It is able to do this using a variety of transport modes including Direct-SAN mode which has the ability to hook straight into the storage layer, thus removing any associated workload from the production network.

It provides a number of options for implementing incremental strategies and leverages Changed Block Tracking (CBT) to do this. The approach does, however, differ in the resulting product of the backup operation. Where a storage snapshotting approach produces a set of intrinsically linked snapshots on primary storage, Veeam Backup & Replication creates a set of deduplicated and compressed backup files on a secondary storage device. These files are just like standard Windows files, which makes them very portable, easy to segment and capable of being moved independently from other backup chain segments.

Veeam Backup & Replication is storage-agnostic and is able to store these files on a wide variety of repository devices. This would include anything that can be made visible to a Windows OS or mounted to a Linux OS. This presents a level of hardware independence that cannot be ignored. You can mix SAN vendors and choose devices more appropriate for storing backups. In addition, you can use much more cost-effective secondary storage (e.g. SATA devices) to further reduce costs.

Where backup data corruption may be a concern, other than basic integrity tests, SANs don't have a mechanism for identifying corruption in snapshots. Corruption would only become visible during the restore process when it fails. Veeam Backup & Replication provides a feature called SureBackup® to provide assurance that backups aren't corrupt and that data is 100 percent restorable. SureBackup doesn't require additional testing hardware. It's able to fulfill its function using minimal host resources and the secondary storage devices used to store the backups. Furthermore, the whole process is automated and requires no manual intervention.

### Granularity issues

Veeam Backup & Replication processes backups in a single pass at the VM level and is able to restore full VMs, individual virtual disks, files and application items from that single backup. Managing the process at the VM level provides much more granularity and provides the ability for the design of the virtual environment to be driven by needs of the business, as opposed to changing it to match the limitations of the data protection process.

### Granular and agentless application consistency

By addressing application consistency at the VM level, you can be much more granular in how and when these tasks are processed. You can invoke application consistency and hypervisor-level snapshots on a per-VM basis rather than in large LUN-level groups. This means you can spread the activity over longer periods, avoiding flash points and optimizing the process to match the capabilities of your host and storage hardware. Some software solutions such as Veeam Backup & Replication can do this without persistent agents, therefore removing any inherent management and resource issues that may be encountered.

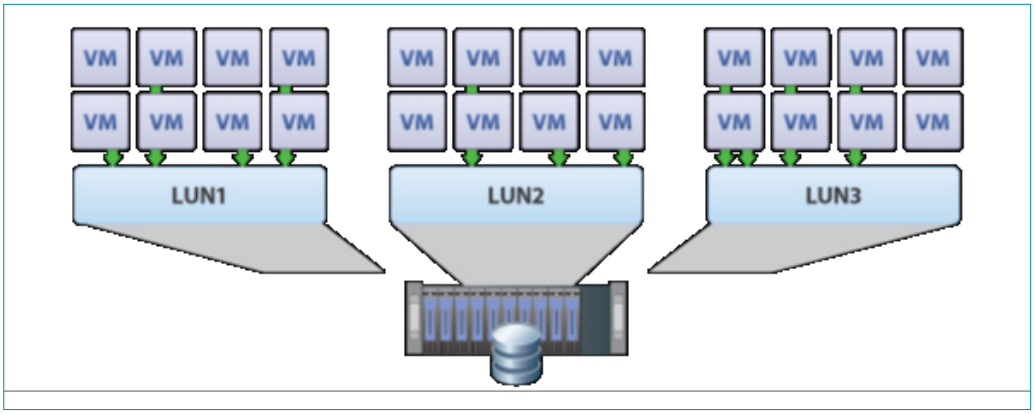


Such an approach would match itself to the virtual environment that is presented, rather than requiring reconfiguration of the environment to support the limitations of the SAN.

### Disk space considerations

By dealing with each backup task at the much more granular VM level, you can also optimize backup storage rather than grabbing and storing bulk LUN-level data. In addition, hypervisors tend to track changes to individual VM disks at a much smaller block size with CBT. This means that there are no massive data increases for relatively small changes (e.g. 15MB increase for a 10KB change).

During the backup process, Veeam Backup & Replication automatically removes data which wouldn't be required in the event of a restore. Empty disk blocks are removed and blocks from in-guest page files are also automatically

removed. When this data has been filtered out, the remaining blocks are then deduplicated and compressed before transmission to the target backup repository. Finally, there is another pass of deduplication at the repository which drastically increases space efficiency.

### Granular restores

Veeam Backup & Replication's restore capabilities are underpinned by block-level access to the data stored in the deduplicated and compressed backup files held at the repository. A very easy-to-use GUI is included to which access can be delegated to first-line team members as well as senior backup administrators. There is also an extremely advanced level of automation provided during restores. In use, these features facilitate the simple interrogation of backup data and subsequent restoration of all levels of data (VMs, virtual disks, files and application items).

Examples of the features supporting these restoration processes include:

- Instant VM Recovery (vPower®): Recover any VM regardless of size in minutes instead of hours.
- 1-Click Restore: Restore any file back to its original location with a single click.
- Veeam Explorer™ for Exchange: Restore any Exchange item back to its original location in a matter of minutes from an agentless backup.
- Veeam Explorer for SharePoint: Search and restore SharePoint items in mere minutes from an agentless backup.

By leveraging the block-level access and advanced automation, these features can not only perform restores without requiring significant additional staging areas, but also provide these restores in five minute timeframes all from a GUI-driven console. Compared with the snapshot recovery approach, RTOs can be significantly reduced when using Veeam. No administrator wants to write command-line scripts at 2 a.m. on a Saturday morning in the event of a failure.

### Hardware dependence, moving backups off-device and long-term retention

As described above, Veeam Backup & Replication is specifically designed for virtualized environments as a storage-agnostic solution, optimized for moving data off the primary production device. Please refer to the previous sections for more details.

Due to the deduplicated, compressed and portable nature of Veeam backup files, both tape and cloud are viable options for storing backups where long-

term retention is required. Veeam Backup & Replication Cloud Edition facilitates this with native tape support introduced with Veeam Backup & Replication v7.

### Backups from storage snapshots

The natural choice is to take the limitations of using storage systems exclusively for the data protection strategy leveraging the good performance of storage snapshots. With Veeam Backup & Replication v7, Veeam introduced the ability to take a VMware backup from a storage snapshot. This builds on the Veeam Explorer for Storage Snapshots feature introduced in version 6.5 that was a recovery technique to leverage storage snapshots for instant VM recovery, file-level recovery and Exchange and SharePoint recovery scenarios.

This is a natural evolution to deliver more performance for your VM backup strategy. Veeam's Backup from Storage Snapshots feature allows VMware VMs to be backed up with the data transfer step being performed from the storage snapshot of HP StoreServ and StoreVirtual products. The benefit of having all of the data transfer performed from the storage snapshot is that the I/O pressure is a mere slice of what it would be otherwise. Furthermore, the VM backup can still be taken with proper application preparation and consistency. Veeam's Backup from Storage Snapshots feature is shown in the figure below:

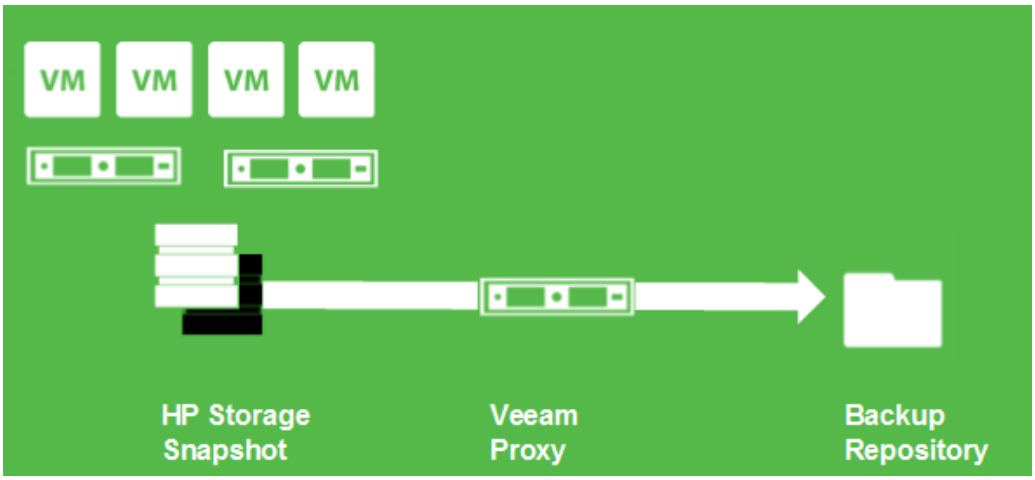


The real takeaway in this approach is that all of the requirements can be met: proper application preparation, high-performance storage snapshots, separation in case of storage failure, deduplication and more—all from an agentless backup.

# Conclusion

### Can I use storage snapshots as backups?

Given the overlapping nature of some of the capabilities we have examined above, it would be very easy to turn this assessment into a storage snapshot backup vs. disk-to-disk backup discussion. In fact, it's not difficult to browse the internet and find the two opposing camps in this debate. The question ultimately being asked is: "Can I use storage snapshots as backups?" One camp says "yes," while the other camp says "no."

The comparison in this paper clearly shows that there are strengths and weaknesses in either approach. The increased speed of the storage snapshotting process and the ability to run very short, incremental cycles to reduce RPOs should be acknowledged as a significant potential benefit. Conversely, the weaknesses in the process related to corruption risks, space requirements, application consistency, granular data management, increased costs and moving data off the primary device cannot be ignored. All of these deficiencies are already resolved in a disk-to-disk backup solution such as Veeam Backup & Replication.

For over a decade the "yes" camp has heralded storage snapshots as being the end of disk-to-disk backups, yet more than a decade later, some of the key proponents of the approach recognize that only a very small percentage of organizations use storage snapshots as backups. Furthermore, they don't expect this balance to tip in the short or medium term, in some cases citing 15-20 year timeframes for storage snapshotting to begin to outnumber disk-to-disk approaches. Where the deficiencies of storage snapshots cause organizations to fail to meet their operational requirements, the same "yes" camp proponents advocate the use of disk-to-disk solutions. From a vendor perspective, this is seen in their technical alliances, with several SAN vendors partnering with and providing integration with disk-to-disk based solutions.

Ultimately, we can conclude that the storage snapshotting capabilities available today and the strategic alliances chosen by SAN vendors give us the answer to our question, and that answer is "No, a snapshot is not a backup."

But, is this the right question to be asking?

## A holistic data protection strategy

Perhaps organizations shouldn't be asking whether to use storage snapshots or disk-to-disk based solutions to protect their data. A more appropriate question might be: "How can I leverage both in a holistic data protection strategy?" Where disk-to-disk solutions may cover all the mandatory requirements of a modern backup solution, storage snapshots could be used alongside them to decrease RPOs and RTOs for recovering data with a short-term retention period.

Using multiple approaches aligned to their most appropriate retention periods provides a powerful and effective strategy for modern data protection. We'll break this down below.

## Short-term retention: Storage snapshot benefits

Storage snapshots clearly provide the fastest method and lowest load for freezing "on-device" data. This means we can take more point-in-time captures of production data and more importantly, do this frequently, perhaps every 30-60 minutes. Doing this with a disk-to-disk solution would increase the load and therefore use less frequent captures in most cases.

In a holistic strategy we would use storage snapshots to do frequent, file-consistent captures of data to circumvent any limitations around granular application consistency. These snapshots could be used for simple LUN and file restores. If it is integrated with a capability such as Veeam Explorer for Storage Snapshots, we could also restore VMs, files and application items in a much simpler, more effective and completely GUI-driven fashion. This would significantly reduce RTOs and be a solid solution for "on-device" protection.

## Medium-term retention: Disk-to-disk benefits

Disk-to-disk backups provide more comprehensive capabilities around application consistency, granular data management, testing restorability, hardware independence and moving data off-device.

Using disk-to-disk backups on a daily or bi-daily schedule would ensure consistent, verified and protected "off-device" backups. This significantly reduces business risk, as well as covering all the necessary criteria for meeting regulatory and compliance requirements. This is important because regulations are becoming more stringent, and most currently require not only that backup strategies are implemented, but also that backups and restores are tested thoroughly.

## Long-term retention: Tape (and/or) cloud benefits

Finally, beyond the day-to-day needs of backing up and restoring data to the production environment, there is also a business need. In most cases, a regulatory requirement to retain backups for long periods is needed in some instances up to seven or ten years. The need to access and restore from these backups is much less frequent and therefore they essentially need to go into long term storage.

The extremely portable, deduplicated and compressed backup files created by Veeam Backup & Replication make them very easy to copy to tape or synchronize with cloud storage.

## Fitting it all together

A holistic strategy would use the right tool for the right job. The use of snapshots, backups, tape and cloud storage all could be considered optimal and most effective, depending on the retention period being targeted.

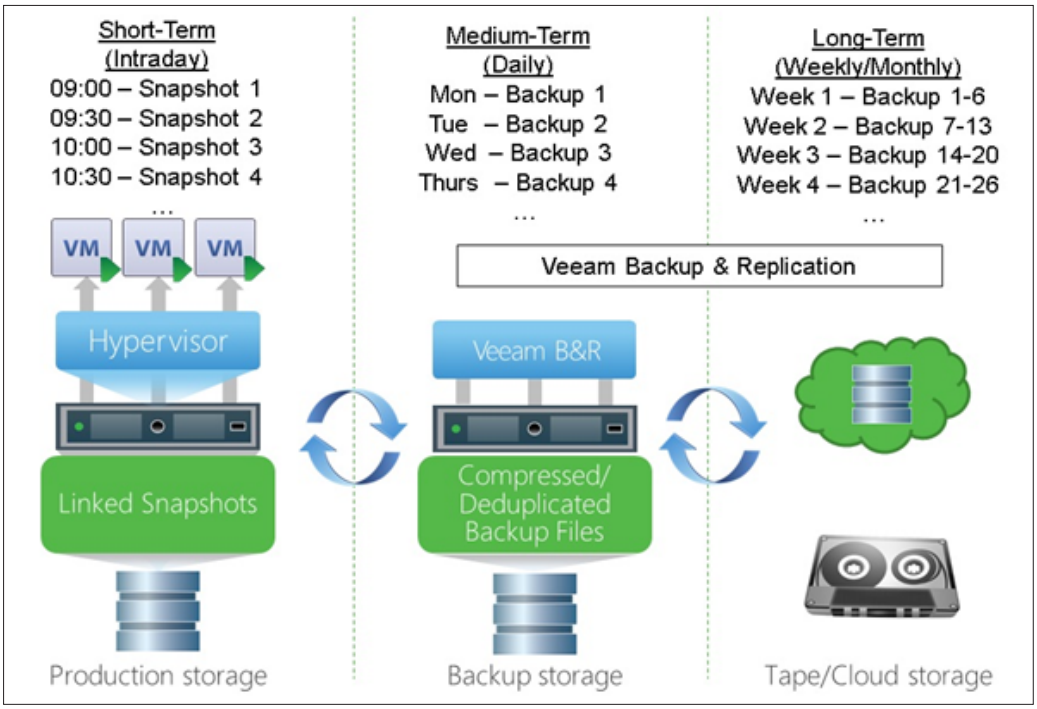


In conclusion, a snapshot is not a backup and it certainly can't fulfill all of the functional and regulatory compliance requirements of a backup solution. It can, however, be considered a powerful data protection tool for modern virtual environments. Used in conjunction with a virtual backup solution such as Veeam Backup & Replication and tape or cloud storage, it can be an important cog in a wider and holistic data protection strategy.

# About the Author

**Hani El-Qasem** is a Senior System Engineer at Veeam Software. He has worked with a broad range of corporate, government and defense intelligence organizations to provide management, data protection and disaster recovery solutions for more than 15 years.

# About Veeam Software

Veeam® is Modern Data Protection™. We believe today's IT requirements have changed and that "3C" legacy backup problems—high costs, increased complexity and missing capabilities—are no longer acceptable for any organization. Veeam provides powerful, easy-to-use and affordable solutions that are Built for Virtualization™ and the cloud—a perfect fit for the modern datacenter.

Veeam Backup & Replication™ is VMware backup, Hyper-V backup, recovery and replication. This #1 VM Backup™ solution helps organizations meet RPOs and RTOs, save time, eliminate risks and dramatically reduce capital and operational costs. Veeam Backup Management Suite™ combines Veeam Backup & Replication and Veeam ONE™ in a single integrated solution to protect virtualization investments, increase administrator productivity and help mitigate daily management risks. Veeam Management Pack™ (MP) extends enterprise monitoring to VMware through Microsoft System Center. Veeam also provides free tools for the virtualization community.

Learn more by visiting http://www.veeam.com.

Veeam
Modern Data Protection

vmware PARTNER
TECHNOLOGY ALLIANCE

Microsoft Partner
Gold Application Development
Gold Management and Virtualization

# Modern Data Protection
## Built for Virtualization

Powerful | Easy-to-Use | Affordable

## Veeam Backup & Replication

## #1 VM Backup for VMware and Hyper-V

Virtualization changes everything – especially backup. If you've virtualized on **VMware or Hyper-V**, now is the time to move up to the data protection solution Built for Virtualization: **Veeam Backup & Replication**.

Unlike traditional backup that suffers from the **"3C" problem** (missing capabilities, complexity and cost), Veeam is:

- **Powerful:** Restore an entire virtual machine (VM) or an individual file, email or database record in 2 minutes
- **Easy-to-Use:** It just works!
- **Affordable:** No agents to license or maintain, works with your existing storage, and includes deduplication, VM replication, Microsoft Exchange recovery, and more!

Join the 58,000 organizations who have already modernized their data protection with Veeam. **Download Veeam Backup & Replication** today!

BEST OF vmworld 2010 GOLD AWARD NEW TECHNOLOGY

BEST OF vmworld 2011 GOLD AWARD NEW TECHNOLOGY

InfoWorld.com ★2013 AWARDS★ TECHNOLOGY OF THE YEAR

Windows IT Pro 2012 Community Choice GOLD

Windows IT Pro 2012 Community Choice SILVER

WINNER 2011 VIRTUALIZATION REVIEW

To learn more, visit **http://www.veeam.com/backup**