

# Cisco HyperFlex Systems and Veeam Backup and Replication

Best practices for version 9.5 update 3 on Microsoft Hyper-V

## What you will learn

This document outlines best practices for deploying Veeam backup and replication software with Cisco HyperFlex™ systems, and the reasoning behind the guidelines for implementation.

Cisco HyperFlex™  
Systems with Intel®  
Xeon® Scalable  
processors



## Contents

|   |           |
|---|-----------|
| Cisco HyperFlex systems and Veeam Backup and Replication 9.5 .....                              | 3         |
| <b>Cisco HyperFlex systems</b> .....  | <b>3</b>  |
| <b>Best practices for deploying Veeam Backup and Replication 9.5 on HyperFlex systems</b> ..... | <b>4</b>  |
| High availability and failover .....  | 4         |
| Data optimization .....   | 4         |
| Data protection .....   | 5         |
| Production checkpoint .....   | 5         |
| Clones .....  | 6         |
| Configuring additional data stores.....   | 7         |
| Virtual machine configuration.....  | 8         |
| Selecting the LUN layout .....  | 8         |
| Selecting the globally unique identifier partition table and file system .....                  | 8         |
| Selecting the allocation unit size .....  | 9         |
| <b>Microsoft Hyper-V configuration</b> .....  | <b>9</b>  |
| Storage quality of service (QoS) .....  | 9         |
| Quality of service .....  | 9         |
| Clustered virtual machine role.....   | 9         |
| <b>Key features in Veeam Backup and Replication 9.5</b> .....                                   | <b>10</b> |
| Veeam backup server .....   | 10        |
| Resource scheduling .....   | 10        |
| Veeam backup change-block tracking .....  | 10        |
| Veeam proxy.....  | 10        |
| Off-host proxy .....  | 10        |
| Backup repository .....   | 11        |
| Network traffic management .....  | 11        |
| Preferred networks .....  | 11        |
| Replication .....   | 11        |
| Instant VM recovery .....   | 12        |
| On-demand sandbox.....  | 12        |
| <b>Conclusion</b> .....   | <b>12</b> |

## Cisco HyperFlex systems

Cisco HyperFlex™ systems use a software-defined infrastructure approach to unlock the full potential of hyperconvergence. They combine software-defined computing in the form of Cisco Unified Computing System™ (Cisco UCS®) servers, software-defined storage with the powerful Cisco® HyperFlex HX Data Platform, and software-defined networking with the Cisco UCS fabric, which integrates with the Cisco® Application Centric Infrastructure (Cisco ACI™) solution.

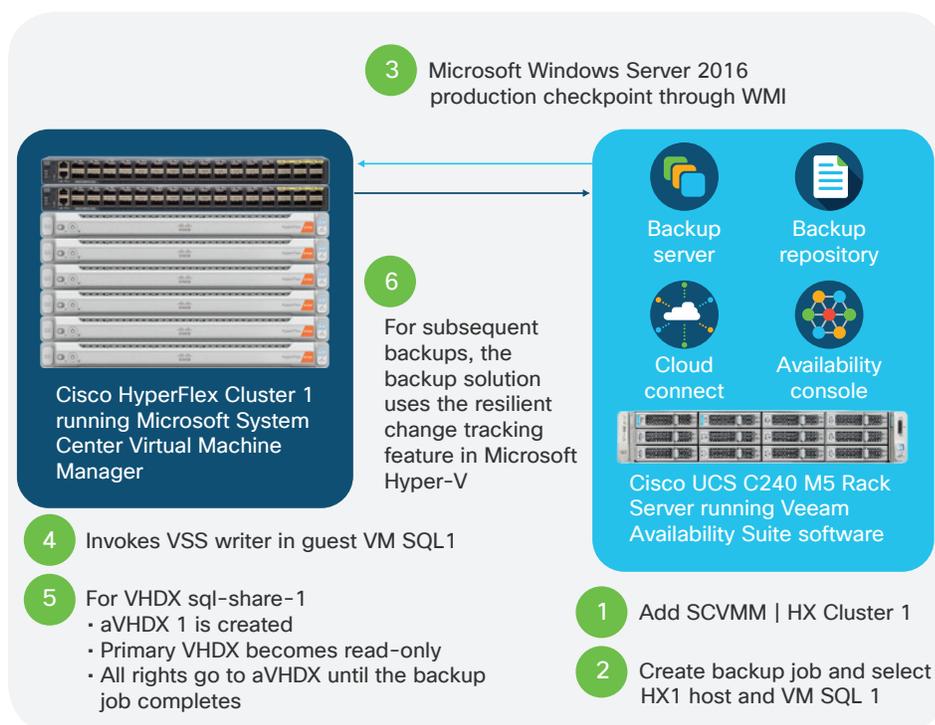
With both hybrid and all-flash memory storage configurations and a choice of management tools, Cisco HyperFlex systems deliver a pre-integrated cluster that is up and running in less than an hour and that scales resources independently to closely match your solution requirements. For an in-depth look at the Cisco HyperFlex architecture, see the Cisco white paper [Deliver Hyperconvergence with a Next-Generation Platform](#).

### Cisco HyperFlex systems and Veeam Backup and Replication 9.5

Veeam Backup and Replication 9.5 creates a production checkpoint on virtual machines (VMs) in a backup job after the Microsoft System Center Virtual Machine Manager (SCVMM) server or individual Microsoft Hyper-V hosts are added to Veeam (Figure 1).

The backup repository can be local to the HyperFlex cluster (e.g., on a Cisco UCS C240 M5 Rack Server), or in a remote data center. You can deploy both a local and a remote copy. To do so, you configure a local backup job, and then create a new backup copy job to copy to the remote office using “From backup files” as the data source.

If you require a large backup repository to support large cluster sizes or multiple clusters, you can use the capacity in Cisco UCS S3260 Storage Servers, which support up to 720 TB of capacity.



**Figure 1.** Cisco HyperFlex systems and Veeam Backup and Replication 9.5 work together to protect data.

## Microsoft Windows Server 2016 and Microsoft Hyper-V

- **Full support.** Cisco HyperFlex systems fully support Microsoft Hyper-V on Microsoft Windows Server 2016.
- **SMB3 file shares.** Data stores are exposed as SMB3 file shares to the hypervisor.
- **Seamless scalability.** Global capacity, compute nodes, and converged nodes scale easily. Nodes and disks are simply added and automatically configured into the system. Data is rebalanced across all available devices in the cluster, increasing the cluster's aggregate capacity and performance.

## Best practices for deploying Veeam Backup and Replication 9.5 on HyperFlex systems

The following sections discuss major design guidelines and configurations for deploying Veeam Backup and Replication 9.5 on Cisco HyperFlex systems. In addition, follow the [best practices](#) outlined by Veeam.

### High availability and failover

High availability is important to messaging administrators. The HX Data Platform in Cisco HyperFlex systems builds in availability at the storage file system layer. All data is written in duplicate or triplicate (replication factor of 2 or 3). A copy of the data can be promoted to primary status if the primary storage controller that owns the primary copy is unavailable, without affecting virtual machines (VMs).

You can configure VMs to be clustered resources in the Failover Cluster Manager or PowerShell software to protect against node or hypervisor failures. The ReadyClone Powershell script provided by Cisco supports the fast cloning of VMs, and the `-AddToCluster` parameter automatically enables clones on HyperFlex systems. You can use the Failover Cluster Manager or a PowerShell command to manually add the cluster machine role to VMs deployed on HyperFlex systems. In the event of a failure or other node outage, VMs are automatically failed over to other nodes in the cluster.

#### PowerShell example:

```
Get-VM <VMName> | add-clustervirtualmachinerole
```

Note that Cisco HyperFlex software upgrades require all non-clustered VMs to be powered off prior to starting the upgrade process. VMs that do not have the cluster VM role do not automatically failover to another node in the Microsoft Windows Failover Cluster.

### Data optimization

The Cisco HyperFlex HX Data Platform includes several data optimization features.

- **Data deduplication** is used on all storage in the cluster, including memory, solid-state disks (SSDs), and in the case of hybrid clusters, hard-disk drives (HDD). By fingerprinting and indexing just these frequently used blocks, high rates of deduplication can be achieved with only a small amount of memory, which is a high-value resource in cluster nodes. Although deduplication rates vary with different workloads, you can experience high rates of large deduplication with operating system and application binary files.

## The advantage of thin provisioning

- **Space savings.** The Cisco HyperFlex HX Data Platform eliminates the need to forecast, purchase, and install disk capacity that may remain unused for a long time.
- **Presented versus actual space allocation.** Virtual data containers (data stores and Virtual Hard Disk [VHDX] files) can present large amounts of logical space to applications, whereas the amount of physical storage space that is needed is determined by the data that is written.
- **Easy expansion.** You can expand storage on existing nodes by adding persistent disks and expand your cluster by adding more storage-intensive converged nodes as your business requirements dictate, eliminating the need to purchase large amounts of storage before you need it. The HX Data Platform automatically expands available cluster storage when new nodes or devices are added.
- **Simplicity.** To increase the size of the thin-provisioned data store, simply edit it using Cisco HX Connect software.

- **High-performance, inline compression** is used on data sets to save storage capacity without negatively affecting performance. Incoming modifications are compressed and written to a new location, and the existing (old) data is marked for deletion, unless the data needs to be retained in a snapshot. The data that is being modified does not need to be read prior to the write operation, which avoids typical read-modify-write penalties and significantly improves write performance.

## Data protection

The solution includes several data protection features.

### Production checkpoint

The Windows 2016 production checkpoint feature can be used to create a VM snapshot. A production checkpoint is a point-in-time image of the VM using backup technology inside the guest to create the checkpoint. This is in contrast with standard checkpoints that save the state, data, and hardware configuration of the VM, which is not suited for production.

Particularly important for application administrators is the capability to take an agentless Microsoft Volume Shadow Copy Service (VSS) application-aware quiesced snapshot. Veeam Explorers can provide recovery of specific individual items. Veeam can take advantage of production checkpoints only if the VM configuration version is 8 or later.

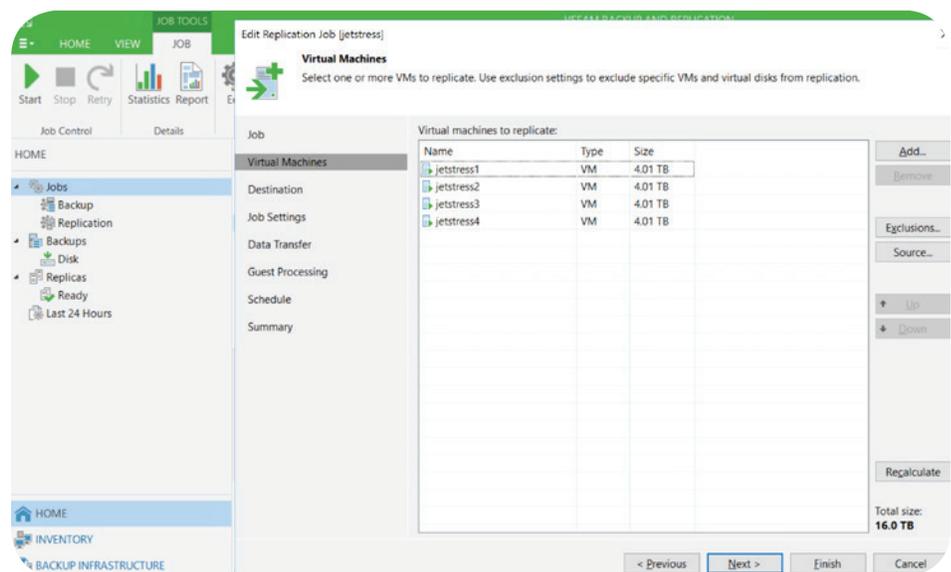


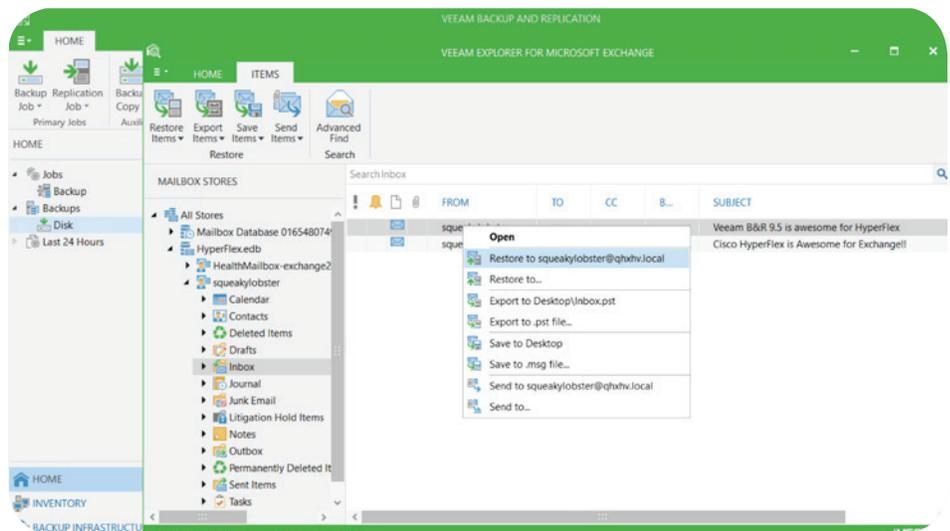
Figure 2. Creating a replication job for VMs

## Benefits of clones

Compared to full-copy methods, ReadyClones can help you:

- Save a significant amount of time
- Increase IT agility
- Improve IT productivity

When a backup or replication job is created, it can contain more than one VM. Each VHDX file in the VM is considered a separate stream when the proxy calculates the number of simultaneous jobs to run. The default for a proxy for “Max concurrent tasks” is equal to the number of CPU cores assigned to the proxy. Each of the same Jetstress VMs has five VHDX files, resulting in 20 tasks to be run. This number is throttled depending on the number of cores associated with the Microsoft Hyper-V host, as it performs as an on-host proxy. Logical cores are excluded, and only physical cores are counted.



**Figure 3.** Restoring individual mail items from a backup

Veeam Explorers, which are custom to each supported application, allow for granular options when mounting an application-aware VSS backup. For example, you can restore individual Microsoft Exchange databases, mailboxes, messages, or other items stored in the database such as contacts (requires a Veeam Enterprise or higher license). These items can be restored to their original location on the production database, to another mailbox, a personal storage table (PST) file, or even a file on the desktop.

## Clones

In the HX Data Platform, ReadyClones can be used to rapidly provision copies of VMs. These fast, space-efficient clones are created through space-efficient metadata operations. Disk space is consumed in the clones only when data is written or changed in the VM. With this approach, dozens of clones can be created and deleted in minutes. Your administrators can easily clone a production VM and test for compatibility with a patch or update. They can then easily remove the clones after testing is complete.

ReadyClones are created on the same data store as the source VM. Moving a VM between data stores is I/O-intensive, requiring the VM to be fully read from the source data store and written to the destination data store.

To provision a ReadyClone, run the HyperFlex ReadyClone PowerShell script with the following four parameters. This example creates three clones of the VM named Veeam, adds the VMs as cluster resources, and names the clones VeeamProxy1, VeeamProxy2, and VeeamProxy3.

```
Run the HyperFlex ReadyClone PowerShell script:  
.\HxClone-HyperV-v3.0.1d-29754.ps1 -VmName Veeam -ClonePrefix VeeamProxy  
-CloneCount 3 -AddToCluster $true
```

ReadyClones are the best-practice way of creating clones on HyperFlex systems. However, the process creates exact copies of VMs, and depending on the application to be installed, may require that you run the `sysprep` command. Read the [“How to Sysprep”](#) blog for more information and step-by-step instructions.

Veeam provides an on-demand sandbox feature which can clone VMs from a backup and place them in an isolated network. This alternative to ReadyClones consumes more storage and runs the VMs from the backup repository, which could have performance implications as backup repositories are often stored on just-a-bunch-of-disks (JBOD) deployments. However, sandbox networking means you can take a group of VMs and stand up their backup images as VMs in a network sandbox. This can be useful, for example, when attempting to patch a Microsoft SharePoint farm with 20 VMs.

### Configuring additional data stores

For most deployments, a single HX Data Platform data store is sufficient, resulting in fewer objects to manage. The HX Data Platform is a distributed file system that is not vulnerable to many of the problems that face traditional systems that require data locality. For example, a VM does not have to fit within the available storage of the physical node. If the cluster has enough space to hold the configured number of copies of the data, the VM will fit. Similarly, moving a VM to a different node in the cluster is a host migration; the data itself is not moved.

In some cases, however, additional data stores may be beneficial. For example, an administrator may want to create an additional HX Data Platform data store to separate a critical database from other workloads. Because performance metrics can be filtered to the data store level, isolation of workloads or VMs may be desired. It is important that all VHDX files within a VM reside on the same data store. The data store is always thinly provisioned on the cluster. The maximum size of the data store is set during data-store creation and can be used to keep a workload, a set of VMs, or end users from running out of disk space on the entire cluster, thereby affecting other VMs.

## VHDX layout

Small VMs run efficiently with a single VHDX, the C: drive, that contains everything. As the number of VM input/output operations per second (IOPS) scales up, isolating each workload onto its own VHDX can increase performance and is a best practice.

Create a separate VHDX for each of the following:

- Operating system
- Paging file
- Databases
- Transaction logs

Each VHDX file is a separate task that is backed up in an independent job through the Veeam proxy, which helps improve backup performance.

## Virtual machine configuration

This section presents some best practices for configuring VMs.

### Selecting the LUN layout

In a VM, the logical unit number (LUN), or logical disk, is a VHDX file stored in the Cisco HyperFlex data store. Usually, after you add a disk to a Microsoft Windows VM, the disk is offline and must be brought online, initialized, and formatted prior to use. Keep in mind that the Veeam software backs up each VHDX file on a separate thread, and the Veeam proxy is limited by the number of simultaneous threads allowed in the configuration.

### Selecting the globally unique identifier partition table and file system

When New Technology File System (NTFS) partitions are initialized, the globally unique identifier (GUID) partition table (GPT) is preferred, as it has more file system redundancy in place and can be used for partitions larger than 2 TB. NTFS has been the recommended and default file system since Microsoft Exchange Server was launched in the 1990s.

Although the new Resilient File System (ReFS) can be used with integrity streams to create and maintain a checksum, doing so can affect performance. Because the HyperFlex HX Data Platform uses checksums at the storage layer, this feature is redundant. The best practice is to format with integrity streams set to false.

Veeam backup repositories can take advantage of fast cloning technology based on block cloning, and should be formatted with ReFS. When deciding to place a Veeam repository on a HyperFlex data store, it is important to ensure no single point of failure. It is recommended that you not place any VMs that reside on a particular HyperFlex cluster into a Veeam backup repository that resides on the same HyperFlex cluster.

**PowerShell example for formatting a disk using ReFS:**

```
Format-Volume -DriveLetter Z -FileSystem ReFS -AllocationUnitSize 65536  
-SetIntegrityStreams $false
```

## Selecting the allocation unit size

For formatting, Microsoft recommends a 64-KB allocation unit size (AUS), sometimes referred to as the cluster size. If you use the default setting, the size will be 4 KB until you have very large partitions. A 64-KB AUS is recommended for disks that house user databases and user transaction logs.

Veeam recommends using a 64-KB AUS for the Veeam repository that will hold all backup files. See the Veeam Backup and Replication best practices document for [repository planning](#) for more information.

```
Format a disk using ReFS and set the AUS to 64 KB:  
Format-Volume -DriveLetter Z -FileSystem ReFS -AllocationUnitSize 65536  
-SetIntegrityStreams $false
```

## Microsoft Hyper-V configuration

This section presents best practices for configuring Microsoft Hyper-V.

### Storage quality of service (QoS)

HyperFlex data stores are presented as SMB3 file shares. As a result, Microsoft failover clustering features do not have any attached storage that can be configured for Microsoft Storage QoS.

### Quality of service

Each VHDX file can be configured with QoS to prevent a VM with runaway IOPS from consuming all cluster resources. To enable QoS on an individual VHDX file, run the following PowerShell command. Be aware that this is in 8 KB I/O sizes; 1 MB I/O is considered as 128 x 8-KB I/O.

```
Enable QoS on a VHDX file:  
Set-VMHardDiskDrive -VMname Veeam -Path "\\HyperFlex.Cisco.Local\DS1\Veeam\  
Veeam1.vhdx" -ControllerType SCSI -MaximumIOPS 2000
```

### Clustered virtual machine role

When a new VM is deployed on a HyperFlex system, it is important to promote the VM to be a clustered VM. Doing so allows the VM to failover to another node in the cluster in the event of planned maintenance or an unplanned outage.

```
Promote a VM to be a clustered VM:  
Get-VM <VMName> | add-clustervirtualmachinerole
```

## Considerations

- **Storage latency control.** In this scenario, the I/O can be throttled when the configured latency threshold is exceeded. This capability is not supported on Cisco HyperFlex systems, because the SMB3 data store is not running on a Windows Server.
- **Concurrent tasks.** Limitations on concurrent tasks can be set on both the backup proxy and repository. Keep in mind each disk within a VM is configured as a separate task. Veeam recommends one core per task be configured at either the proxy or repository layer.
- **Backup repositories.** These repositories can be limited in read and write data rates. This can help you to size the rate to match what the backup repository storage can handle and indirectly limit the throughput affecting the source VM storage. If latency impact to the HyperFlex storage is unacceptable during backup, attempt to limit the number of simultaneous tasks at the proxy and repository layer, and limit the read and write data rates on the backup repository.

## Key features in Veeam Backup and Replication 9.5

### Veeam backup server

The Veeam backup server is the Microsoft Windows Server on which the Veeam Backup and Replication v9.5 update 3 software is installed. It coordinates backup, replication, and restore tasks and serves as a default backup proxy. The exception is when Microsoft Hyper-V is used. In this case, the hosts serve as on-host proxies.

### Resource scheduling

Veeam provides two VM data processing modes; Parallel and Sequential

- **Parallel processing.** This mode allows backup, replication, and restore jobs to process multiple virtual disks and VMs simultaneously up to the limits set on the backup infrastructure.
- **Sequential processing.** This mode instructs the Veeam software to process VM disks and VMs in a backup job sequentially. If “enable parallel processing” is disabled, Veeam infrastructure components can still run more than one job in parallel up to the limits set on the backup infrastructure.

### Veeam backup change-block tracking

During incremental jobs, Veeam data movers use change-block tracking to copy on the blocks that have changed since the last backup. Microsoft calls this resilient change tracking (RCT), and it persists even if the VM is live migrated to a different HyperFlex host.

### Veeam proxy

During a backup or replication job, the Microsoft Hyper-V host (where the VM resides) is used as the proxy by default. During on-host backup, the Microsoft VSS components and the Veeam Data Mover are invoked on the source host. In the on-host backup mode, the Veeam software reads from VM disks in read-only mode and removes the production checkpoint after moving the changed blocks to the backup repository. Removing the production checkpoint causes the avhdx files to be merged with their original vhdx files.

### Off-host proxy

Off-host proxies use transportable shadow copies, and are not supported on HyperFlex systems, because it is not supported to connect Microsoft Hyper-V servers outside of the HyperFlex cluster to the HyperFlex SMB3 data store.

### Backup repository

A backup proxy, or folder on the backup storage, is used to store backup chains and metadata for replicated VMs. A backup repository can be hosted on a Microsoft Windows Server with block storage, a Linux server with block storage, or an a network file system (NFS) share, an SMB share, and a deduplicating appliance. It is not supported to place a Veeam backup repository directly on the HyperFlex SMB3 data store.

### Network traffic management

Many basic backup applications read entire data sets or all changed blocks since the last backup, with data transfers completing as fast as the storage system and operating system can process them. Because Cisco HyperFlex systems are built on Cisco UCS, with fast 10 Gigabit Ethernet ports on each host, this process can result in multiple gigabytes per second of backup throughput with just a few simultaneous backup jobs. You can throttle a pair of IP address ranges on both the source and target components between which data is transferred over the network.

### Preferred networks

Preferred networks can be configured for backup and replication traffic by setting the network address using Classless Inter-Domain Routing (CIDR) notation. A best practice is to configure the live migration network to be separate from other networks so that it can be specified in the Veeam configuration and to provision the Veeam repository server to be on that virtual LAN (VLAN).

### Replication

The Veeam Backup and Replication replicates VMs. A replication job makes an exact copy of the VM on the target Microsoft Hyper-V host and keeps it up to date according to the replication interval.

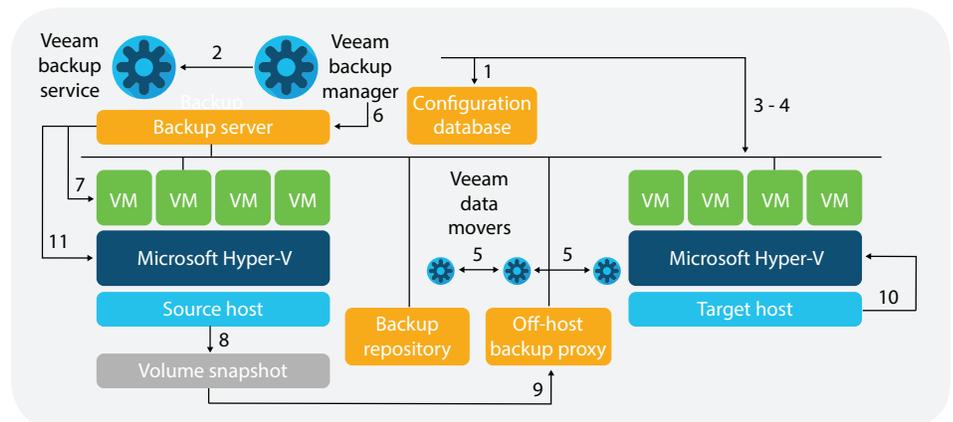


Figure 4. The replication process

## For more information

- [Cisco HyperFlex systems](#)
- [Deliver Hyperconvergence with a Next-Generation Data Platform](#) white paper
- [Hyper-V Backup and Storage Best Practices](#), a Veeam white paper

### Instant VM recovery

Instant VM recovery allows you to restore a VM into your production environment by running it directly from the compressed and deduplicated backup file in the Veeam repository. If the Veeam repository is not on sufficiently performing storage, consider restoring the backup to either the original or an alternate location located on HyperFlex storage.

### On-demand sandbox

The on-demand sandbox feature uses instant VM recovery mechanisms to start one or more VMs from a backup to perform tasks such as troubleshooting VM problems, testing patches, and installing new software. The on-demand sandbox uses a virtual lab in an isolated environment.

## Conclusion

The Cisco HyperFlex HX Data Platform revolutionizes data storage for hyperconverged infrastructure deployments that support new IT consumption models. The platform's architecture and software-defined storage approach provides a purpose-built, high-performance distributed file system with a wide array of enterprise-class data management services. With innovations that redefine distributed storage technology, the data platform gives you the hyperconverged infrastructure you need to deliver adaptive IT infrastructure.

Cisco HyperFlex systems in hybrid and all-flash configurations lower both operating expenses (OpEx) and capital expenditures (CapEx) by allowing you to scale as you grow. They also simplify the convergence of computing, storage, and network resources. Size and acquire what you need now, and easily scale your storage with automated rebalancing after you add disks to the converged nodes or add more converged nodes. If more computing resources are required, use both Cisco UCS approved rack and blade servers, adding them to the cluster as compute-only nodes.

Unlike many traditional storage systems, you can easily increase the size of the VHDX or your data store. Native cloning is space-efficient and fast, and provides your HyperFlex administrators with quick access to production data for testing and optimization without requiring additional storage capacity. And running Veeam software on Cisco HyperFlex systems provides additional recovery functions with Microsoft VSS application-consistent snapshots and replication between Microsoft Hyper-V hosts.