

HyperFlex Edge Site Disaster Recovery with Veeam Backup & Replication

Contents

Prerequisites	3
Cisco HyperFlex	3
Veeam	3
Introduction	3
Solution overview	3
Configuration overview	4
HyperFlex Edge site	4
Core data-center Veeam Backup & Replication server and HyperFlex recovery cluster	4
Terminology	4
Backup infrastructure	5
Storage infrastructure	5
Job types	5
Failover and failback	6
Step-by-step configuration	6
Administration	15
Failover	15
Documentation references	16
Document information	16

Prerequisites

Cisco HyperFlex

We recommend reviewing the Cisco HyperFlex® HX Data Platform release notes, installation guide, and user guide before proceeding with any configuration. The data platform should be installed and functioning as described in the installation guide. Please contact Cisco® Technical Assistance Center (Cisco TAC) or your Cisco representative if you need assistance.

Veeam

An understanding of the deployment and use of Veeam Backup & Replication is recommended. Additionally, Veeam documentation is available at: <https://www.veeam.com/documentation-guides-datasheets.html?productId=8&version=product%3A8%2F221>.

Introduction

The distributed nature of edge-site architectures introduces several challenges related to data protection and disaster recovery. Performing local backups with the ability to conduct local recovery operations is one requirement. Another challenge involves edge-site disaster recovery. Planning for inevitable edge-site outages, be they temporary, elongated, or permanent, requires a disaster recovery solution. This document provides information about orchestrating DR failover and failback for Cisco HyperFlex Edge cluster disaster recovery with Veeam Backup & Replication.

Solution overview

A number of well-known disaster recovery topologies include the active/active, active/standby, and many-to-one solutions. From a cost perspective, a many-to-one solution, with centralized data-center resources as a disaster recovery facility for multiple edge sites, represents a financially viable solution when compared to active/active and active/standby solutions. The solution presented in this document is focused on the many-to-one topology. Conceptually, a number of HyperFlex Edge clusters utilize a single HyperFlex data-center cluster as a disaster recovery resource. Veeam Backup & Replication provides data management capabilities for backups, offsite backup copies, and replication of virtual machines (VMs).




Figure 1.
Many-to-one replication example

Configuration overview

Veeam Backup & Replication is installed in the core data center. Required Veeam Backup & Replication services are push-installed to a VM on each HyperFlex Edge site. The required Veeam Backup & Replication services typically include backup transport, vPower NFS, and mount service. Each HyperFlex Edge site will also require Veeam Backup & Replication backup infrastructure components consisting of a backup proxy and backup repository. Veeam Backup & Replication HyperFlex Edge site deployments can be hosted on one or more virtual machines, or on a Cisco UCS host (for example, a Cisco UCS C240 rack server).

At a high level, backups are performed locally on each HyperFlex Edge cluster. Backups are retained in the local repository for fast and flexible recovery (for example, granular, instant, or full recovery). The local HyperFlex Edge backups are then automatically copied from the local edge repository to the core site repository. At the point where backups have been copied to the core repository, they are then automatically replicated to a datastore on the HyperFlex recovery cluster. In effect, there are three retained copies of the VM data, a backup in the HyperFlex Edge site repository to fulfill local recovery objectives, a backup in the core site repository to fulfill offsite recovery objectives, and a VM in the HyperFlex recovery cluster datastore to fulfill failover objectives. This process satisfies the 3-2-1 backup rule. It is important to note that different retention values can be set for each location; local, core, and datastore. It is also important to note that Veeam Backup & Replication orchestrated failover is distinctly different from the flexible recovery (instant, granular, or full) of VMs from backup. Veeam Backup & Replication failover occurs on VM replicas that already reside in a production datastore and are registered within VMware. As a result, Veeam Backup & Replication failover operations are near-instantaneous and do not require recovering VM data into a datastore or a VMware instant recovery operation that utilizes storage vMotion. Veeam Backup & Replication enables one-click orchestrated failover with capabilities to automatically re-IP VM replicas and network re-mapping on the HyperFlex recovery cluster.

Network requirements consist of connectivity between HyperFlex Edge sites and a core data center. The overall design is largely a customer choice based on business requirements. Bandwidth requirements of the replication network should be understood based on protected workloads and data change rates.

HyperFlex Edge site

A HyperFlex Edge site is typically a remote or branch office, a retail location, or manufacturing location. The HyperFlex cluster is deployed as two, three, or four edge-specific nodes.

Core data-center Veeam Backup & Replication server and HyperFlex recovery cluster

The core data center includes disaster recovery resources consisting of an on-premises Veeam Backup & Replication deployment on Cisco UCS servers, and a HyperFlex cluster with the role of hosting recovered HX Edge site workloads. The core site HyperFlex cluster can also be used to host production VMs if it has been properly sized for such usage.

Terminology

The terms used by various data management, backup, and disaster recovery products varies widely. This section takes a look at Veeam terminology that may be different when compared to other products. In particular, the terms “backup copy” and “replication” as used by Veeam should be understood because they may differ from other products.

Backup infrastructure

Veeam Backup & Replication backup proxies and backup repositories are both considered components of a Veeam backup infrastructure.

A Veeam VMware backup proxy resides between the backup server and other components of the backup infrastructure. The backup proxy processes jobs and delivers backup traffic. This solution requires a backup proxy on each HyperFlex Edge site as well as at the core data center. Additional information about Veeam Backup & Replication backup proxies is available here:

https://helpcenter.veeam.com/docs/backup/vsphere/backup_proxy.html?ver=120.

A Veeam Backup & Replication backup repository is a component that manages a storage location for retaining backup files, backup copies, and metadata for replicated VMs. This solution requires a backup repository on each HyperFlex Edge site as well as at the core data center. Additional information about Veeam backup repositories is available here:

https://helpcenter.veeam.com/docs/backup/vsphere/backup_repository.html?ver=120.

Storage infrastructure

Within the context of this solution, the Cisco HyperFlex clusters are all considered components of a Veeam Backup & Replication storage infrastructure. Configuring the HyperFlex clusters as storage infrastructure enables Veeam Backup & Replication to leverage HyperFlex native snapshot technology. This functionality may be referred to as storage integration. This integration enables backups from storage snapshots to perform edge backups in the fastest and most efficient manner. Additional information about HyperFlex requirements and limitations for the integrated solution is available here:

https://helpcenter.veeam.com/docs/backup/vsphere/storage_limitations_cisohx.html?ver=120.

Job types

This solution employs the use of three specific Veeam Backup & Replication job types; backup, backup copy, and replication. To avoid confusion, the purpose and result of each different job type is discussed below.

- Backup – Veeam Backup & Replication backup jobs produce image-level backups of VMs. They treat VMs as objects, not as a set of files. When you back up VMs, Veeam Backup & Replication copies a VM image as a whole, at a block level. Additional information is available here: <https://helpcenter.veeam.com/docs/backup/vsphere/backup.html?ver=120>.
- Backup copy – Veeam Backup & Replication backup copy jobs facilitate the creation of an additional instance of the same backup file and copy each instance to a secondary (target) backup repository. Target backup repositories can be located in the same site as the source backup repository or can be deployed off-site. The backup copy file has the same format as the primary backup. Restores/recoveries can be performed directly from a target repository in the event of a disaster. Note that other vendors and products may refer to this functionality as “replication.” Additional information is available here: https://helpcenter.veeam.com/docs/backup/vsphere/about_backup_copy.html?ver=120.
- Replication – Veeam Backup & Replication jobs create exact copies of the VMs in the native VMware vSphere format on a target host VMware datastore and register powered-off VMs with VMware vCenter. Veeam Backup & Replication maintains these copies in sync with each source VM. Replication can provide minimum Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs) in case a disaster strikes because VM replicas reside in an in a ready-to-start state. Additional information is available here: <https://helpcenter.veeam.com/docs/backup/vsphere/replication.html?ver=120>.

Failover and failback

Failover is the process of switching from the VM on the source HyperFlex Edge site to its VM replica on a target HyperFlex cluster in the core data center. Failover works in conjunction with replica VMs that are created by means of replication jobs. Replica VMs reside in a VMware datastore and are registered within VMware. From a simplistic perspective, a failover operation powers on the replica VMs and executes very quickly.

Failback is a process of returning production operation from the VM replica to the source VM. The failback operation consists of two phases. During the first phase, Veeam Backup & Replication synchronizes the state of the production VM (the source VM, an already recovered VM, or a VM that will be recovered from the replica) with the current state of its active replica. The second phase switches all processes from the active VM replica to the production VM, turns off the replica, and also sends to the production VM changes made to the VM replica since the end of the first phase.

Additional information is available here:


https://helpcenter.veeam.com/docs/backup/vsphere/failover_failback.html?ver=120.

Step-by-step configuration

The steps provided assist in configuring existing core-data-center Veeam Backup & Replication deployments to provide local-edge-data protection, remote backup copies, and replication to build out a HyperFlex Edge site data-protection and disaster-recovery solution.

1. HyperFlex Edge site and core-data-center backup proxies

In the Veeam Backup & Replication user interface, select “Backup Infrastructure” and then select “Backup Proxies.” The solution calls for a VMware backup proxy located at each HyperFlex Edge site as well as at the core data center.



Backup Infrastructure		Type in an object name to search for		
		Name	Type ↓	Host
Backup Proxies		HX-Edge-London	VMware	HX-Edge-London
Backup Repositories		HX-Edge-NewYork	VMware	HX-Edge-NewYork
External Repositories		HX-Edge-Sydney	VMware	HX-Edge-Sydney
Scale-out Repositories		VMware Backup Proxy	VMware	S3260-SanJose
WAN Accelerators		Backup Proxy	Agent	S3260-SanJose
Service Providers				

Figure 2.

VMware backup proxies for three HyperFlex Edge sites and a core data center

Adding a new VMware backup proxy requires a VM or physical host running a supported operating system deployed at a HyperFlex Edge site or at a core data center. Additional information about deploying a VMware backup proxies is available here:

https://helpcenter.veeam.com/docs/backup/vsphere/backup_proxy.html?ver=120.

2. HyperFlex Edge site backup repository

In the Veeam Backup & Replication user interface, select “Backup Infrastructure” and then select “Backup Repositories.” The solution calls for a backup repository located at each HyperFlex Edge site as well as the core data center.

Backup Infrastructure		Type in an object name to search for						
	Name ↑	Type	Host	Path	Capacity	Free	Used Space	
Backup Proxies								
Backup Repositories	HX-London-Repository	Windows	HX-Edge-London	R:\Backups	1023.9 GB	975.5 GB	42.1 GB	
External Repositories	HX-NewYork-Repository	Windows	HX-Edge-NewY...	R:\Backups	1023.9 GB	971.8 GB	45.8 GB	
Scale-out Repositories	HX-Sydney-Repository	Windows	HX-Edge-Sydney	R:\Backups	1023.9 GB	969.4 GB	48.3 GB	
WAN Accelerators	S3260-Repository	Windows	S3260-SanJose	R:\Backups	1023.9 GB	868.3 GB	149 GB	

Figure 3. Backup repositories for three HyperFlex Edge sites and a core data center

Additional information about deploying a backup repository is available here: https://helpcenter.veeam.com/docs/backup/vsphere/backup_repository.html?ver=120.

3. Storage Infrastructure for HyperFlex Edge site and core data center

In the Veeam Backup & Replication user interface, select “Storage Infrastructure” and, if shown, select “Cisco HyperFlex.” The solution calls for Cisco HyperFlex storage infrastructure at each HyperFlex Edge site as well as at the core data center.

Storage Infrastructure		Name ↑	Description
Storage Infrastructure		HyperFlexLondon	Created by S3260-SANJOSE\Administrator at 5/15/2023 4:44 PM.
Cisco HyperFlex		HyperFlexNewYork	Created by S3260-SANJOSE\Administrator at 5/15/2023 4:45 PM.
		HyperFlexSanJose	Created by S3260-SANJOSE\Administrator at 5/15/2023 4:58 PM.
		HyperFlexSydney	Created by S3260-SANJOSE\Administrator at 5/15/2023 4:51 PM.

Figure 4. Cisco HyperFlex clusters configured as storage infrastructure at three HyperFlex Edge sites and a core data center

Adding a Cisco HyperFlex cluster to Veeam Backup & Replication as storage infrastructure is a straightforward process. The workflow involves specifying the HyperFlex DNS name or IP address, admin user credentials, and which backup proxy (or proxies) to use. A recommended prerequisite step to be executed on the backup proxy is to add a NIC (network interface connection) that is connected to the HyperFlex hypervisor data network. A static IP address is recommended.

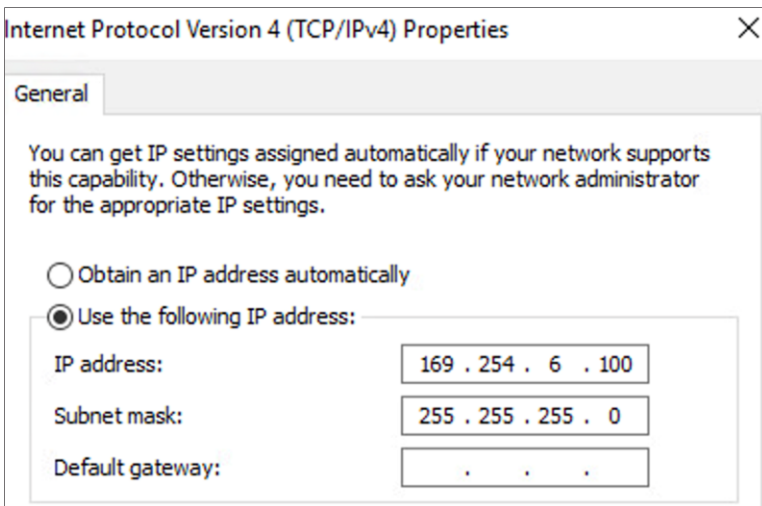


Figure 5. NIC on the HyperFlex hypervisor data network configured with a static IP address

The assigned hypervisor data network IP address should be configured as an “Allowed IP address” for the ESXi NFSAccess service within the firewall settings on each ESXi node of the HyperFlex cluster. Note that the use of the ESXi firewall “Allow connections from any IP address” setting is not recommended because it may introduce a security risk.

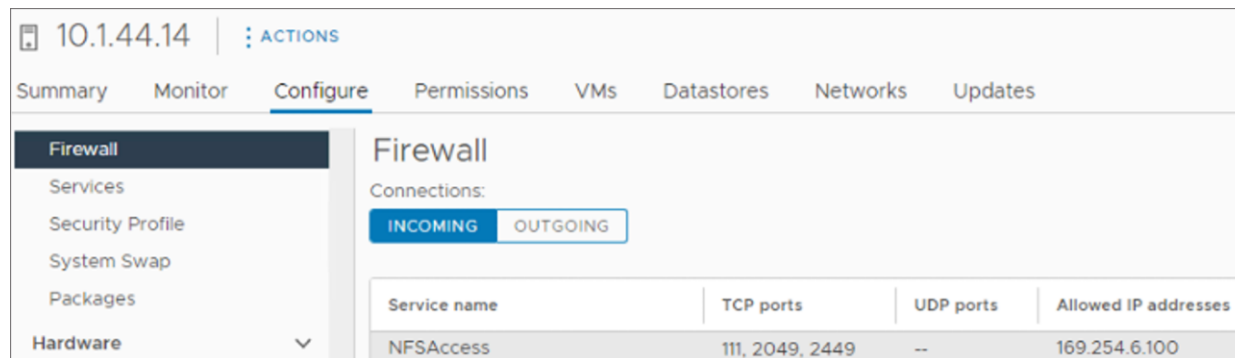


Figure 6. ESXi host firewall configuration allowing NFSAccess to the static IP address assigned to a VMware backup proxy

These important steps enable the use of the Veeam Backup & Replication D-NFS (direct NFS) transport mode when performing backup and restore operations. Scanning a HyperFlex cluster configured as storage infrastructure will indicate a successful completion when the static IP address on the HyperFlex hypervisor network has been properly configured.

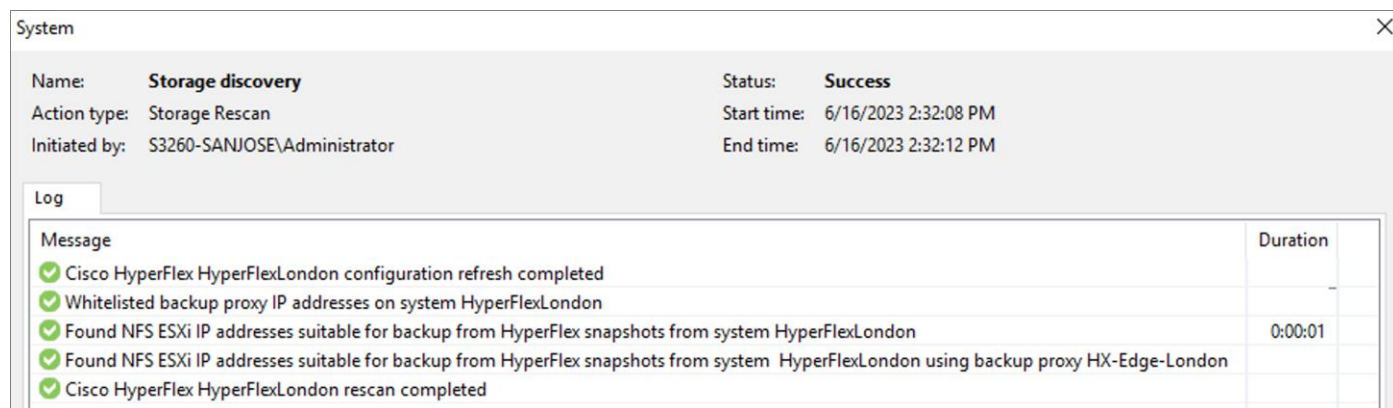


Figure 7. Storage rescan results of a HyperFlex Edge cluster configured as storage infrastructure named “HyperFlexLondon” configured with a VMware backup proxy named “HX-Edge-London”

Additional information about adding Cisco HyperFlex clusters as storage infrastructure is available here: https://helpcenter.veeam.com/docs/backup/vsphere/cisco_backup_configure.html?ver=120.

4. Backup jobs

Include HX Edge VMs that require protection when selecting the VMs to back up. The backup proxy selection should be the local backup proxy in the same HX Edge site as the VMs that are being protected. Similarly, the backup repository selection should be the local backup repository in the same HX Edge site as the VMs being protected. Select a retention policy that meets business requirements for local backups. The retention period can be specified in “days” or in “restore points.” Note that the number of offsite backup copies retained, and the number of replicas retained, are configured independent of retention parameters used for local backups.

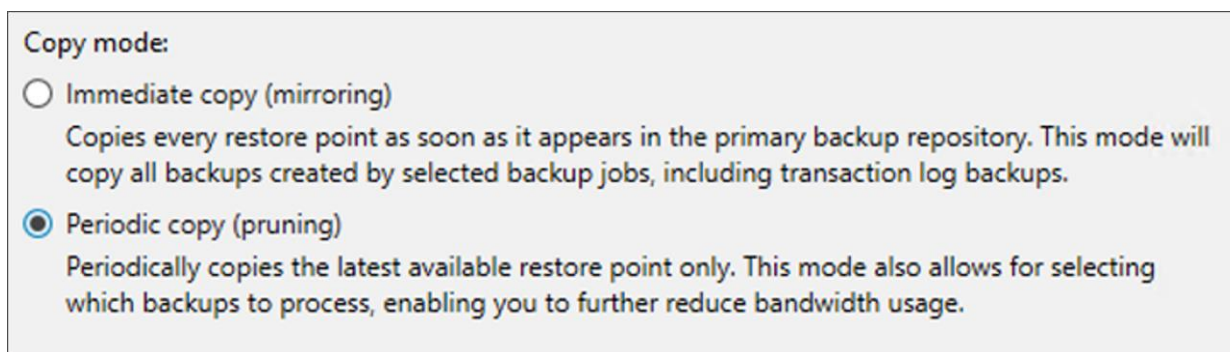
Enable the check-box named “Configure secondary destinations for this job.” If a backup copy job has not yet been created, clicking the “Next” button will allow the job configuration process to continue. If a backup copy job has already been created, it can be selected as a secondary destination job within the backup job.

5. Backup copy jobs

The purpose of a backup copy job is to create a copy of HX Edge site VM backups in a different backup repository. This solution utilizes the backup copy job to create a copy of HX Edge site VM backups in a remote backup repository located in the core data center. The backup copies serve as both an off-site copy of VM backups, and as a source for replication jobs (which are discussed in the “Replication jobs” section of this document). The net benefits of the copy job are twofold: first, in the event of an HX Edge site outage, VM backups are available for recovery at the core data center; and, second, the creation of VM replicas that are used for failover are sourced from backup copies, eliminating the need to transfer data from the edge to the core twice.

A backup copy job includes a “copy mode” selection in which either immediate or periodic copies are created. The immediate copy mode uses mirroring to copy every restore point from the source backup repository. The periodic copy mode copies only the latest restore point, enabling a capability to selectively copy specific restore points from the source backup repository. Additional information about the copy mode is available here:

https://helpcenter.veeam.com/docs/backup/vsphere/backup_copy_modes.html?ver=120.



Copy mode:

- Immediate copy (mirroring)
Copies every restore point as soon as it appears in the primary backup repository. This mode will copy all backups created by selected backup jobs, including transaction log backups.
- Periodic copy (pruning)
Periodically copies the latest available restore point only. This mode also allows for selecting which backups to process, enabling you to further reduce bandwidth usage.

Figure 8.

Selection of the “Periodic copy” mode within the “Job” step when configuring a backup copy job

Selecting the workloads to copy to the target backup repository is the next step. An object to be copied can be selected from either a job or a backup repository. As examples, a backup job that protects HX Edge site VMs, or a local backup repository at an HX Edge site, can be selected as the object to process. Additional information about selecting the workload to copy is available here:

https://helpcenter.veeam.com/docs/backup/vsphere/backup_copy_vms.html?ver=120.



Objects to process:			
Name	Type	Size	
 HX-Edge-London-Backup	VMware Backup Job	51.4 GB	<input type="button" value="Add..."/> <input type="button" value="Remove"/>

Figure 9. Selection of a backup job in the “Objects to process” step when configuring a backup copy job

The target backup repository should be a backup repository located within the core data center. The retention policy also needs to be specified.



Target
Specify the target backup repository, number of recent restore points to keep, and the retention policy for full backups. You can use map backup functionality to seed backup files.

Job	Backup repository:
Objects	S3260-Repository (Created by S3260-SANJOSEAdministrator at 5/15/2023 4:20 PM.)
Target	863 GB free of 0.99 TB Map backup
Data Transfer	Retention policy: 14 <input type="button" value="restore points"/> <input type="checkbox"/> Keep certain full backups longer for archival purposes Configure...

Figure 10. Selection of a backup repository located within the core data center. (A retention value of 14 restore points has also been configured.)

The data transfer method, either direct or through built-in WAN accelerators, needs to be selected. WAN accelerators, which perform data reduction / data deduplication of the backup data transferred from edge to core sites, are Veeam Backup & Replication components that are deployed in pairs, with one WAN accelerator deployed at the HX Edge site and the other WAN accelerator deployed at the core data center. Additional information about Veeam WAN accelerators is located here:

https://helpcenter.veeam.com/docs/backup/vsphere/wan_accelerator.html?ver=120.

Figure 11.
Selection of direct transfer mode

Scheduling a backup copy job is the next step in the process. The options presented are based on the copy mode selected earlier in the job creation process. If the immediate copy mode has been selected, the user is prompted to edit a data transfer window that determines when the job can transfer data.

If the periodic copy mode has been selected, the options presented include a check-box, “Run the job automatically,” which should be selected. There are a number of possible scheduling options, including daily job execution, monthly job execution, periodic job execution, or the ability to run the backup copy job immediately after a preceding job.

Figure 12.
“Run the job automatically” and “After this job” scheduling options. (A local HX Edge site backup job has been selected from the pull-down menu. The backup copy job is scheduled to execute immediately following the local HX site backup job.)

There is one additional step, which will configure automatic execution of a replication job. This step will be covered in the subsequent “Replication Jobs” section.

6. Replication jobs

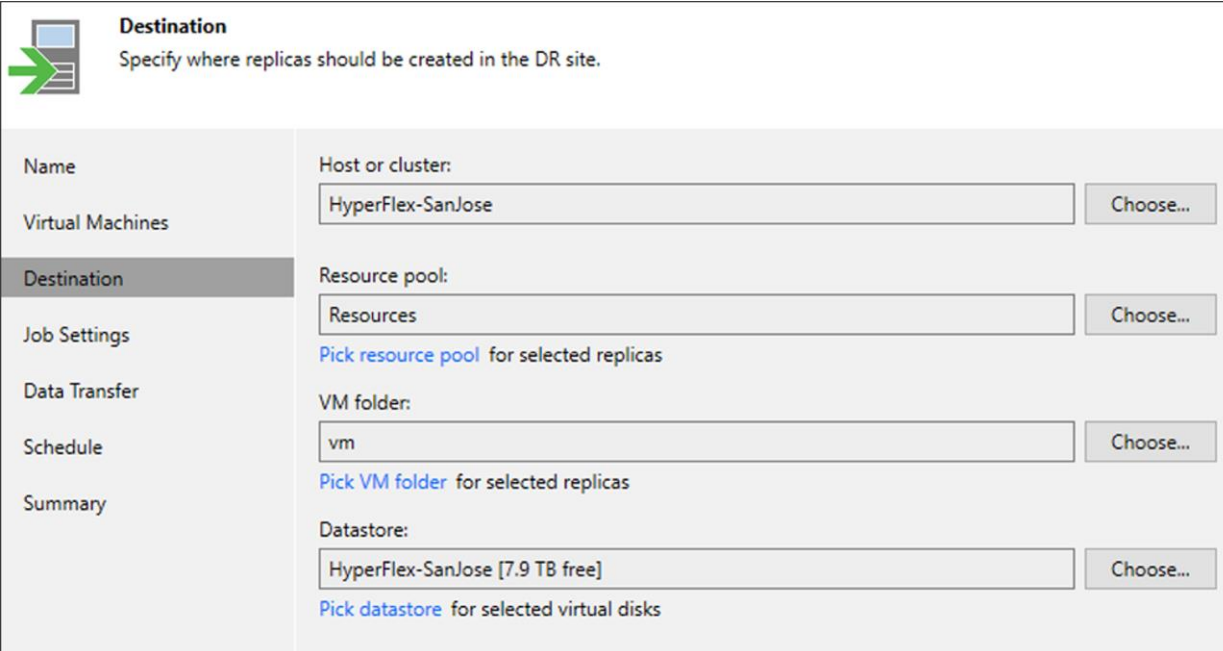
Creation of a replication job begins with inputting a job name followed by optionally selecting up to three advanced controls; these include “Replica seeding,” “Network remapping,” and “Replica re-IP.” The advanced controls are not discussed in this document; however, additional information is available for each topic.

- “Replica seeding” information is available here:
https://helpcenter.veeam.com/docs/backup/vsphere/replica_create_seed.html?ver=120.
- “Network remapping” and “Replica re-IP” information is available here:
https://helpcenter.veeam.com/docs/backup/vsphere/replica_name_vm.html?ver=120.

The purpose of a replication job is to create VM replicas within a target ESXi datastore. In the case of this solution, the target ESXi datastore is hosted by a HyperFlex recovery cluster. The VM replicas are used to perform failover and failback operations, which are discussed in a subsequent section of this document.

Proceed to create a replication job by selecting the VMs that should be replicated. The source should be set to be the backup repository residing within the core data center, located in close proximity to the HyperFlex recovery cluster. Within the selected repository, select the VMs from the HX Edge site that should be replicated.

The next step is to specify where the replicas should be created in the disaster recovery site. This includes the ESXi host or cluster name, which should be the HyperFlex recovery cluster. The resource pool and VM folder can also be specified. The datastore selection should be a datastore residing on the HyperFlex recovery cluster.



The screenshot shows the 'Destination' configuration step in a Veeam replication job. The interface has a sidebar on the left with navigation options: Name, Virtual Machines, Destination (selected), Job Settings, Data Transfer, Schedule, and Summary. The main area is titled 'Destination' and contains the instruction 'Specify where replicas should be created in the DR site.' Below this are four configuration fields, each with a text input and a 'Choose...' button: 'Host or cluster' (HyperFlex-SanJose), 'Resource pool' (Resources), 'VM folder' (vm), and 'Datastore' (HyperFlex-SanJose [7.9 TB free]). There are also blue links for 'Pick resource pool for selected replicas', 'Pick VM folder for selected replicas', and 'Pick datastore for selected virtual disks'.

Figure 13.
Replication destination step of replication

In the Job Settings step of the replication job creation process, select the backup repository for replica metadata. This is the backup repository located in the HX Edge site. If required, the default replica name suffix “_replica” can be altered. Also, specify the number of restore points to keep.

The data transfer step of the replication job creation process should be configured to use the target proxy located in the core data center. There is also an option to select the direct or WAN-accelerated data transfer options. Because the backup repository and target ESXI datastore are both located in the core data center, the “Direct” setting is recommended for this solution.

Within the scheduling step, it is not required to run the job automatically or to select any scheduling parameters. The reason for this is that the solution executes the replication job through a script that gets executed automatically after the backup copy job completes. At this point a PowerShell script should be created that will be used to execute the replication job.

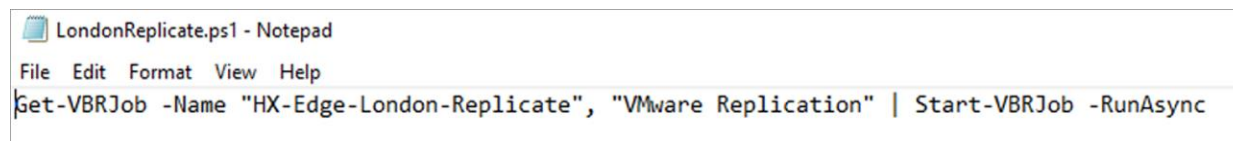
A screenshot of a Notepad window titled "LondonReplicate.ps1 - Notepad". The window has a menu bar with "File", "Edit", "Format", "View", and "Help". The main text area contains a single line of PowerShell code: `Get-VBRJob -Name "HX-Edge-London-Replicate", "VMware Replication" | Start-VBRJob -RunAsync`.

Figure 14.

Text of a PowerShell script that, when invoked, will cause a replication job named “HX-Edge-London-Replicate” to execute

The Veeam PowerShell Reference document is available here:

https://helpcenter.veeam.com/docs/backup/powershell/veeam_psreference.html?ver=120.

Once a PowerShell script has been created, it can be tested manually using the “Veeam Backup & Replication PowerShell Toolkit.” To access the toolkit from the Veeam Backup & Replication user interface, select “Console” > “PowerShell.”

The final configuration step involves editing the backup copy job to invoke the PowerShell script. Within the backup copy job, select the “Target” step and then click the “Advanced” button. Within the “Advanced Settings” window, select the “Scripts” tab. Enabling the check-box “Run the following script after the job” allows browsing for the PowerShell script. After selecting the PowerShell script, radio buttons are available to configure the frequency or cadence of PowerShell script execution.

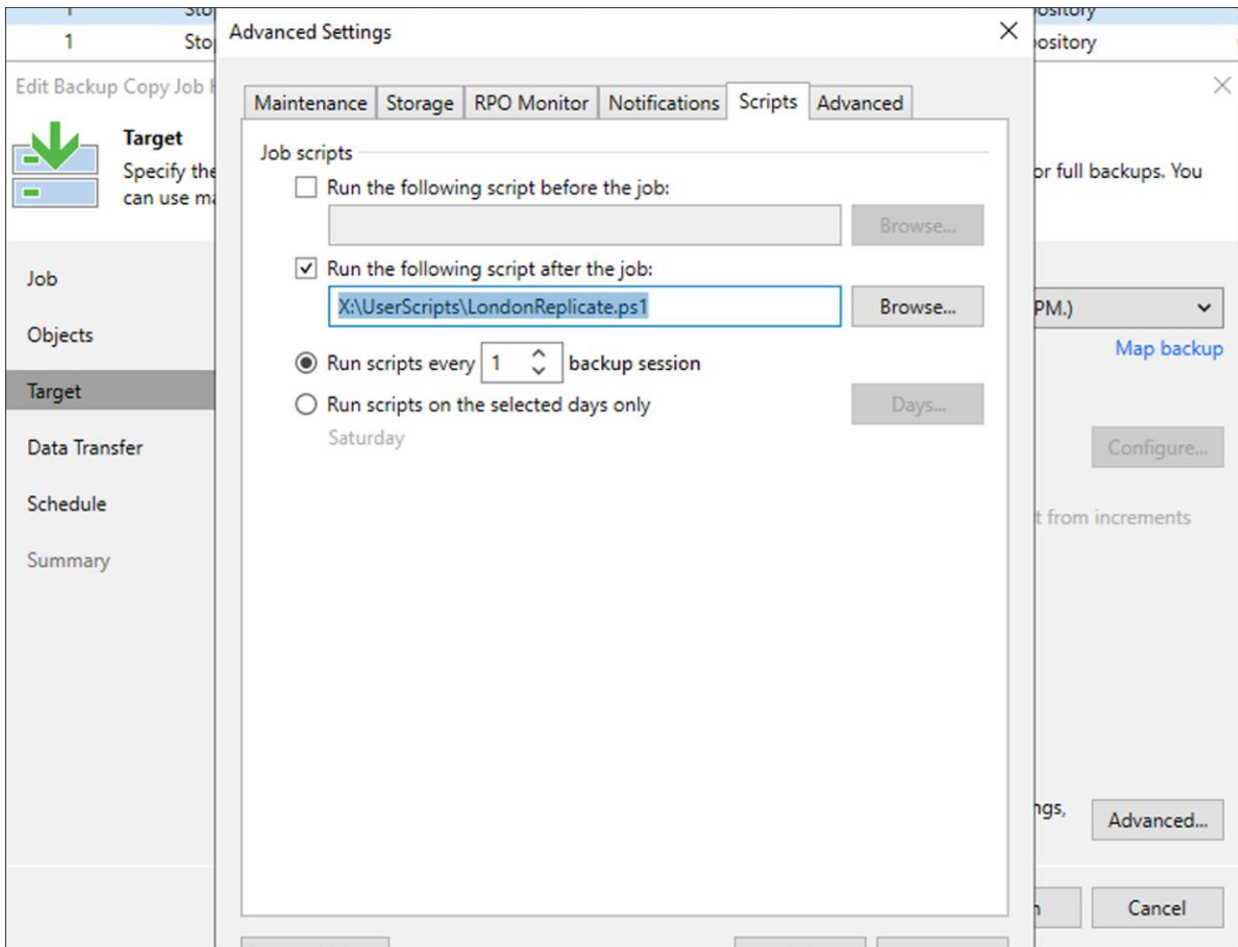


Figure 15. Example of a backup copy job “Target” “Advanced Settings” “Script” window, which can be used to call a PowerShell script

At this stage, the backup, backup copy, and replication jobs have been configured and scheduled to execute without user intervention.

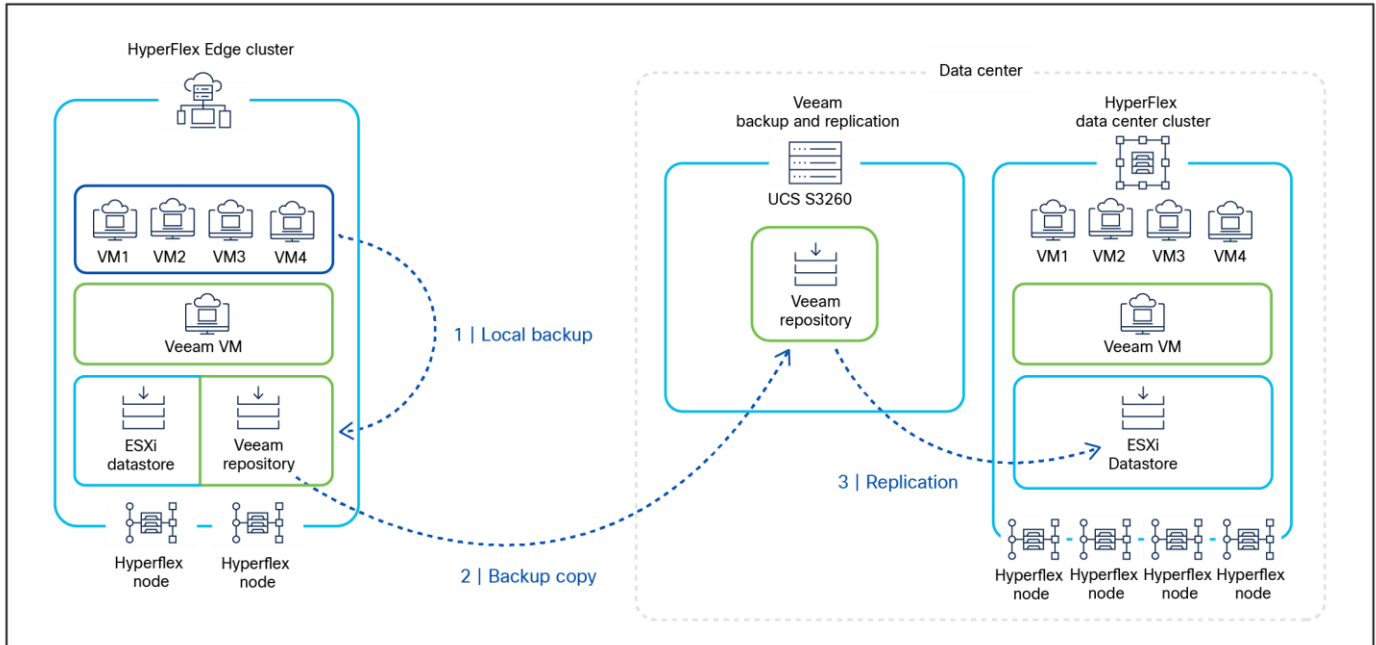


Figure 16.
Automatic execution of backup, backup copy, and replication jobs

Administration

Failover

Failover is a process of switching from the VM on the source host to its VM replica on a target host in the disaster recovery site. Failover is a much faster process than a typical restore operation. The reason for this is that the VM data is already located in a production EXSi datastore and registered within vSphere as the result of a replication job, so the time to failover is effectively the VM boot time.

There are three methods to conduct a failover, and it is essentially a user choice as to which method to select:

- **Manual failover** can be invoked on one or more VMs simultaneously. The process allows the user to select among available recovery points. Manual failover does not allow for the orchestration of the sequence in which VMs are powered on.
- **Failover plans** allow for the orchestration of the power-on sequence of replicated VMs, adding a time-delay value between each VM in a power-on sequence, and the execution of pre- and post-failover scripts.
- **Planned failover** can be executed manually with the added functionality of triggering a replication job to perform an incremental replication job run to get the latest changes to the replica VM. Planned failover also offers the ability to specify a time delay between each VM's power-on sequence.

After a failover is initiated, the process requires a final step referred to as finalization. In effect, finalization completes the failover process. There are three methods to conduct finalization. The choice of which method to use should be based on the expected outcome of the failover.

- Undo failover discards all changes made to replica VMs while in the failover state. Undo failover can be used when testing failover. The undo failover method of finalization will effectively end a test failover session when testing has been completed.
- Failback to production switches back to the production VMs from the replica VMs. When failing back to the original production VM in its original location, Veeam Backup & Replication will transfer only the changes that occurred on the replica VMs during the failover event back to the source/recovered VM. The failback to production method of finalization requires an additional step called “commit failback.” Commit failback is performed when the VM that was failed back (the production VM) works as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.
- Permanent failover permanently switches from the original source VMs to their replica VMs. The replica VMs stop acting as replicas and start acting as production VMs. The original source VMs are excluded from the replication job. The permanent failover method of finalization cannot be undone and cannot failback to production.

Documentation references

Access to the following documents is recommended because they contain content referenced within this paper:

- Cisco HyperFlex HX Data Platform documentation: <https://www.cisco.com/c/en/us/support/hyperconverged-systems/hyperflex-hx-data-platform-software/series.html>.
- Veeam Help Center technical documentation: <https://www.veeam.com/documentation-guides-datasheets.html?productId=8&version=product%3A8%2F221>.

Document information

Document summary	Prepared for	Prepared by
HyperFlex Edge Site Disaster Recovery with Veeam Backup & Replication V1.0	Cisco Field	Bill Roth, Mark Polin

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)