

HyperFlex Edge Site Disaster Recovery with Cohesity

Version 1.0

June 2023

Contents

Prerequisites	3
Cisco HyperFlex	3
Cohesity	3
Introduction	3
Solution overview	3
Configuration overview	4
HyperFlex Edge sites	4
Central data center Cohesity clusters and HyperFlex recovery clusters	4
Terminology	4
Storage domains	4
Remote clusters	5
Policies	6
Sources	6
Protection groups	7
Step-by-step configuration	7
Administration	9
Failover	9
Cloning	10
Documentation references	10
Document information	11

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 HX Data Platform should be installed and functioning as described in the installation guide. Please contact Cisco® Technical Assistance Center (TAC) support or your Cisco representative if you need assistance.

Cohesity

An understanding of the deployment and use of Cohesity DataProtect is recommended. Additionally, the ability to log in to the Cohesity web portal is also recommended. The Cohesity web portal provides access to Cohesity product documentation, Cohesity downloads, and Helios.

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 the inevitable edge site outage, be it temporary, elongated, or permanent requires a disaster recovery solution. This document provides information about HyperFlex Edge cluster disaster recovery with Cohesity DataProtect.

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, the use of centralized data-center resources as a disaster-recovery facility for multiple edge sites represents a financially viable solution when compared to the active/active and active/standby solutions. The solution presented in this document is focused on the many-to-one topology. Conceptually, some number of HyperFlex Edge clusters utilize a single HyperFlex data-center cluster as a disaster recovery resource. Cohesity DataProtect provides data-management orchestration for backups, replication, and recovery.

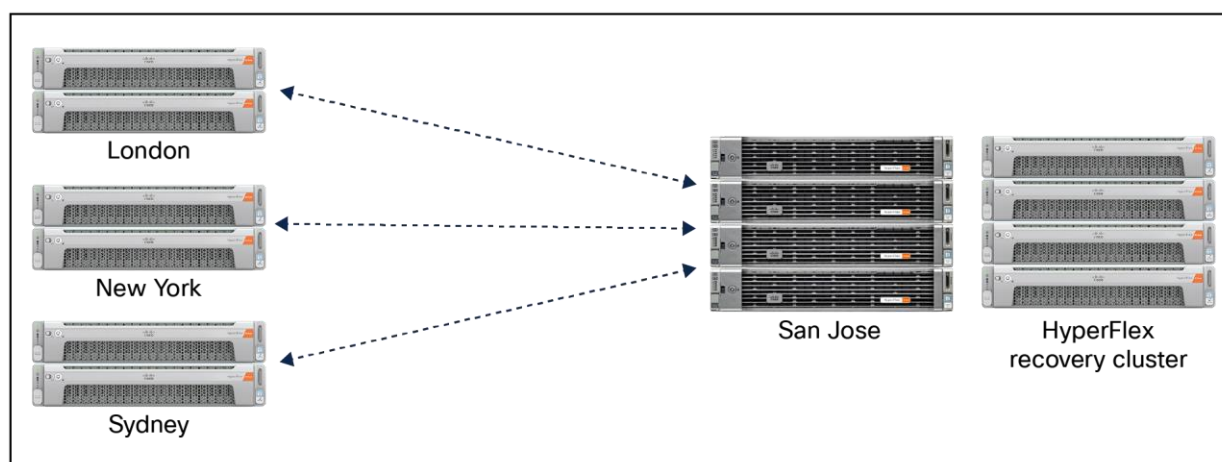


Figure 1.
Many-to-one replication example

Configuration overview

Each HyperFlex Edge site, as well as the central data center, requires its own Cohesity DataProtect deployment. As an example, ten HyperFlex Edge sites would require ten DataProtect deployments, and the central data center would require an additional DataProtect deployment. Each Cohesity DataProtect deployment is referred to as a cluster, regardless of whether the deployment resides on physical nodes, consists of clustered virtual machines, or consists of a single virtual machine.

Replication of protected data occurs between different Cohesity clusters. Replicated data consists of snapshots, which represent recovery points for use in failover events or failover testing (cloning). Network requirements consist of connectivity between connected Cohesity clusters, and 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 sites

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

Cohesity DataProtect is available as a virtual or a ROBO edition. The virtual edition is deployed as a virtual machine directly on a HyperFlex Edge cluster. It can be deployed in a small or large configuration.

The ROBO edition is deployed on a Cisco UCS® C220-M5 Rack Server. The ROBO edition delivers improved performance and increased storage capacity when compared to the virtual edition.

Central data center Cohesity clusters and HyperFlex recovery clusters

The central data center includes disaster recovery resources consisting of an on-premises clustered Cohesity DataProtect deployment on Cisco UCS nodes and a HyperFlex cluster that hosts recovered HyperFlex Edge site workloads.

Terminology

Storage domains

A storage domain is a named storage location on a Cohesity cluster. A storage domain defines the policy for deduplication and other configuration settings such as compression, and encryption. There are also advanced settings for quotas, alerts, and views. Reports can be generated that provide detailed usage statistics for storage domains, including logical-data-protected, storage-consumed, and storage-growth rates.

Storage domain pairing is required for replication. When configuring replication, a storage domain on a source Cohesity cluster is paired with a storage domain on a remote Cohesity cluster. It is important to note that multiple source-storage domains can be paired to a single remote-storage domain when the source-storage domains reside on different Cohesity clusters. For example, storage domains associated with ten HyperFlex Edge sites can be paired with a single storage domain associated with a central data center.

Table 1. Paired storage domains

Storage Domain Pairing	
Local Storage Domain	Remote Storage Domain
London Storage Domain	Sanjose Storage Domain
New York Storage Domain	Sanjose Storage Domain
Sydney Storage Domain	Sanjose Storage Domain

Protection groups specify the use of a specific storage domain for storing snapshots. Protection groups are discussed in greater detail in a subsequent subsection of this document.

Remote clusters

The replication of data from a source Cohesity cluster to a remote Cohesity cluster requires registering a remote cluster. The registration process consists of identifying the remote cluster using an IP address (or addresses) and providing a username and password. By default, the “Auto Select” interface setting is enabled. This setting can be altered to use a specific interface group.




Figure 2.

Three remote Cohesity clusters connected to a central data center Cohesity cluster

Additional cluster options can also be selected during the registration process. The first is the remote access setting, which enables management of the remote cluster from the current cluster. When enabled, the remote cluster can be selected from a drop-down menu, simplifying administrative access to the remote cluster. The second option enables replication. When enabled, storage domain pairing and a number of replication parameters can be configured. The replication parameters include outbound data compression, encryption, data transfer rate limiting, and data transfer rate limit overrides.

Policies

A policy consists of a backup schedule and retention parameters. For example, perform a backup at some frequency and retain the backup for some duration. Options to perform periodic full backups and to apply extended retention rules are also available.

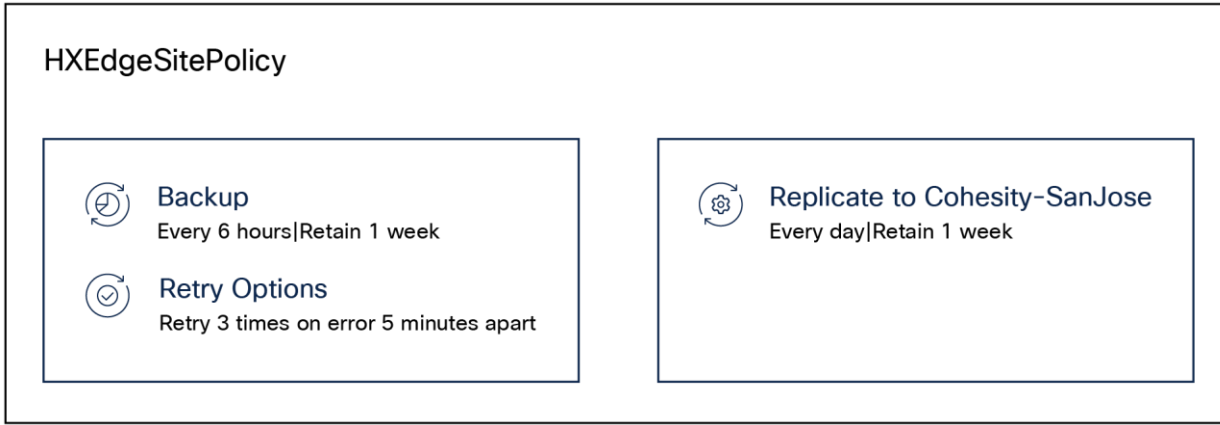


Figure 3.
Example of a policy that includes replication

Replication can be added to a policy. When adding replication to a policy, a remote-cluster name is selected from a list of previously configured remote clusters. A new remote cluster can also be registered when adding replication. Replication parameters include frequency-based scheduling as well as a retention period.

Sources

A source is defined as an object that should be protected, such as a virtual server, storage array, or database. Sources are a registered entity within Cohesity DataProtect. Examples of sources are VMware vCenter hosts and HyperFlex clusters.

The screenshot shows a form titled "Register Storage Snapshot Provider". It includes a dropdown menu for "Storage Snapshot Provider: Hyperflex", a text input for "Hostname or IP Address" with the value "10.10.10.10", and fields for "Username" (admin) and "Password". There are "Register" and "Cancel" buttons at the bottom.

Figure 4.
Registering a HyperFlex cluster as a source

Registering a HyperFlex cluster as a source enables Cohesity DataProtect to leverage HyperFlex native snapshots.

Protection groups

A protection group includes virtual machines to protect, a policy that defines execution, replication, and retention periods, a storage domain, and additional settings. Among the additional settings is the “Leverage Storage Snapshots for Data Protection” property. When this is enabled and “HyperFlex” is selected from the drop-down menu, the protection group will use HyperFlex-native snapshots whenever possible.

Step-by-step configuration

The steps provided assist in configuring existing Cohesity DataProtect deployments to provide local data protection (backups) and disaster recovery (replication). Cohesity documentation for initial cluster setup is available from Cohesity at: <https://docs.cohesity.com/Setup/Welcome/InitialClusterSetup.html>. Note that access to Cohesity documentation requires a “MyCohesity Login” user ID and credentials.

1. Storage domain creation

The default configuration includes a storage domain named “DefaultStorageDomain.” Creating one or more storage domains that align with the organization’s standard naming conventions is recommended. The use case for creating multiple storage domains on source and remote clusters is typically based on accounting and reporting requirements. For instance, when multiple source clusters map their local storage domain to a common shared remote storage domain, space reporting on the remote storage domain will reflect an aggregated set of metrics. The opposite is true when the remote cluster has a separate mapped storage domain for each source cluster. In this scenario, space reporting will be granular and reflect the space consumption of each source cluster individually.

Storage domains are easy to create by selecting “Settings > Summary > Storage Domains > Create Storage Domain.” Detailed Cohesity online documentation for creating or editing storage domains is available [here](#).

2. Add remote cluster

Creating a new cluster connection begins with inputting the IP address of the remote cluster along with a user ID and password for the remote cluster. At the point where the connection is validated, two cluster options are available for selection. The first option is remote access, which can be enabled or disabled based on business requirements. The second option is replication, which needs to be enabled in order to proceed with configuring the disaster recovery solution. The next step is to add storage domain pairing. A local storage domain is selected by means of a pull-down menu, and a remote storage domain is selected by means of a separate pull-down menu. At this point a number of replication settings can be applied, including outbound compression, encryption, and data-transfer rate limiting.

Adding a remote cluster is accomplished by selecting “Infrastructure > Remote Clusters > Add Cluster.” Detailed Cohesity online documentation for adding a remote cluster is available [here](#).

3. Policy creation

The default installation automatically creates a number of preconfigured policies. The “Bronze,” “Silver,” and “Gold” policies are based on predetermined recovery point objectives with backup frequencies of 1 day, 12 hours, and 4 hours and retention periods of 1 month, 2 weeks, and 1 week respectively. Replication can be added to an existing policy, or a new policy can be created.

A policy includes a scheduled backup frequency and backup retention period. Extended retention parameters can be added. Scheduled periodic full backups can also be added to the policy. Adding replication to a policy involves selecting the remote cluster, scheduled replication frequency, and a retention period for replicated backups on the remote cluster.

Policy creation is initiated by selecting “Data Protection > Policies > Create Policy.” Detailed Cohesity online documentation for creating a policy is available [here](#).

4. Registering sources

The sources that need to be registered for data protection and disaster recovery include the hypervisor (VMware vCenter) and HyperFlex cluster. Specifically, these are the vCenter that manages the local HyperFlex cluster and the local HyperFlex cluster itself. The vCenter source is registered so that managed virtual machines can be selected for inclusion within a protection group. The HyperFlex cluster source is registered so that the HyperFlex API can be used to manage HyperFlex native snapshots.

Registering the vCenter source is accomplished by selecting “Data Protection > Sources > Register > Virtual Machines.” Similarly, registering the HyperFlex cluster source is accomplished by selecting “Data Protection > Sources > Register > Storage Snapshot Provider.” Detailed Cohesity online documentation for registering a hypervisor source is available [here](#). Detailed Cohesity online documentation for registering a storage snapshot provider is available [here](#).

5. Creating a protection group

The protection group construct includes objects that require protection, a user-assigned protection group name, a policy, and a storage domain. The objects to be protected are associated with the source vCenter that was already created and will consist of virtual machines. The user-assigned protection group name should be descriptive and align with the organization’s standard naming conventions. The selected policy should be the policy already created that includes replication to a remote cluster. The selected storage domain should be the storage domain already created. The additional setting “Leverage Storage Snapshots for Data Protection” should be enabled and the “HyperFlex” storage snapshot provider should be selected.

Creating a policy is accomplished by selecting “Data Protection > Protection > Protect > Virtual Machines.” Detailed Cohesity online documentation for creating a protection group is available [here](#).

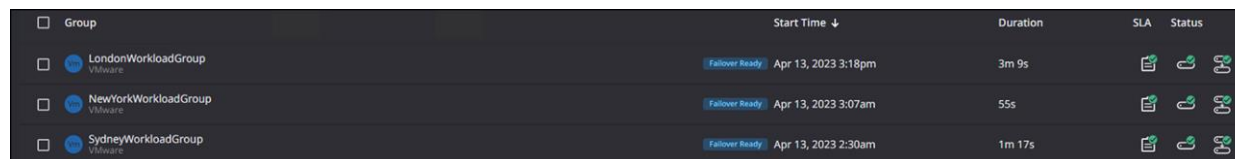
Administration

There are two management interfaces available with which to administer Cohesity DataProtect. The on-premises management interface directly connects to a Cohesity cluster by means of an IP address. This interface can be used for single cluster management. Optionally, remote access to one or more connected clusters can be configured; this enables management of the remote cluster from the current cluster. When enabled, a remote cluster can be selected from a drop-down menu, simplifying administrative access to the remote cluster.

Helios, a SaaS management console, adds to the capabilities of the on-premises interface with multicluster management, support automation, and analysis. Access to Helios requires login credentials. Login credentials are associated with a specific Helios realm. It's important to note that, when adding a Cohesity cluster to Helios, the cluster will be added to the realm the user is assigned to.

Failover

When a protection group with a replication executes, a local backup will be performed followed by replication. After the first snapshot for each VM in the protection group is replicated, a failover-ready copy of the protection group is created on the remote Cohesity cluster. A replicated protection group that is in a failover-ready state will be displayed as "Failover Ready."



Group	Start Time ↓	Duration	SLA	Status
LondonWorkloadGroup VMware	Failover Ready Apr 13, 2023 3:18pm	3m 9s		
NewYorkWorkloadGroup VMware	Failover Ready Apr 13, 2023 3:07am	55s		
SydneyWorkloadGroup VMware	Failover Ready Apr 13, 2023 2:30am	1m 17s		

Figure 5.

Three failover-ready protection groups

Clicking the ellipses (on the action menu) to the right of the protection status will cause a pop-up menu to appear with options to "Failover" or "Delete" the protection group. Failover will cause the protection group to be activated on the Cohesity cluster. Failover will also cause new incoming replicated snapshots to be rejected. The failover workflow facilitates recovery from replicated snapshots retained on the cluster. The recovery process enables recovery of the entire protection group, or a subset of the virtual machines that have been protected. At the point where the selection of the entire protection group or a subset of virtual machines to recover has been made, the user can then select a specific recovery point from which to perform the recovery operation. By default the most recently replicated recovery point is selected. The user can override the default recovery point selection and list available recovery points by means of a pull-down menu and then select any retained recovery point.

A number of options are then presented. A "Recover To" location can be specified, and this should be set to the central data center's HyperFlex recovery cluster. When selecting the recovery location, the resource pool, datastore, and VM folder are also selected. The recovery method defaults to VMware instant recovery. This behavior can be altered to perform a copy recovery if desired. Additional recovery options include network, naming, and power state selections. Detailed Cohesity online documentation about protection group failover is available [here](#).

Additional notes about failover:

- After initiating a protection group failover, incoming replication from the original HyperFlex Edge site protection group is blocked. The original HyperFlex Edge site Cohesity cluster may or may not be operational depending on the disaster that made a failover necessary. The user will need to manually deactivate the original protection group when possible on the original HyperFlex Edge site Cohesity cluster. Detailed Cohesity online documentation for deactivating the original protection group is available [here](#).
- At the point where failover has completed and the original HyperFlex Edge site Cohesity cluster is operational, configuring replication from the central data center's Cohesity cluster back to a HyperFlex Edge site Cohesity cluster is a manual operation.
- The protection group that was used to conduct the failover operation will need to be updated in order to protect the recovered VMs. Edit the protection group to add the virtual machine objects into the protection group. Replication back to a HyperFlex Edge site Cohesity cluster can also be added to the protection group policy as required. Detailed Cohesity online documentation for restoring the protection group to a running state is available [here](#).
- When reverse-direction replication has been configured and successfully executed back to a HyperFlex Edge site Cohesity cluster, another failover operation can be conducted. This will allow the HyperFlex Edge site to resume running production workloads.

Cloning

The ability to test-recover virtual machines is an important part of disaster-recovery planning and preparedness. This action is performed as a clone operation on a Cohesity cluster. Clones can be created on either a source or remote Cohesity cluster. The clone operation creates a copy of a recovery point that is presented in a temporary Cohesity datastore and mounted on the ESXi hosts of a HyperFlex cluster. The clone workflow allows for selection of a retained recovery point, selection of a location for the clones, and the naming of the temporary Cohesity datastore on which the clones will reside. Additional parameters include adding a name prefix/suffix, as well as network and power state selections.

Cloning is accomplished by selecting "Test and Dev > Clone > VMs." Detailed Cohesity online documentation about creating and tearing down clones is available [here](#).

Documentation references

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

- [Cohesity Initial Cluster Setup](#)
- [Cohesity Version 6.6 Documentation](#)
- [Cohesity Downloads](#)

Document information

Document summary	Prepared for	Prepared by
HyperFlex Edge Site Disaster Recovery with Cohesity V1.0	Cisco Field	Bill Roth

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)