



Managing Virtual Machine Disaster Recovery

- [HX Data Protection Snapshot Overview, on page 1](#)
- [Protecting Virtual Machines Overview, on page 8](#)
- [Disaster Recovery Overview, on page 28](#)
- [Replication Maintenance Overview, on page 37](#)

HX Data Protection Snapshot Overview

The HyperFlex Data Protection Snapshot (DP Snapshot) feature enables the protection of virtual machines from disaster by configuring the replication of running VMs between a pair of network connected clusters. Protected virtual machines running on one cluster replicate to the other cluster in the pair, and vice versa. The two paired clusters typically are located at a distance from each other, with each cluster serving as the disaster recovery site for virtual machines running on the other cluster.

Once protection is configured on a VM, the HX Data Platform periodically takes a DP Snapshot of the running VM on the local cluster and replicates (copies) the DP snapshot to the paired remote cluster. In the event of a disaster at the local cluster, the most recently replicated snapshot of each protected VM can be recovered on the remote cluster. Each cluster that serves as a disaster recovery site for another cluster must be sized with adequate spare resources so that upon a disaster, it can run the newly recovered VMs in addition to its normal workload.



Note Only the most recently replicated DP snapshot is retained on the destination cluster. Retaining additional DP snapshots is not supported.

Each VM is individually protected by assigning it protection attributes, including the replication interval (schedule). The shorter the replication interval, the fresher the replicated snapshot data is likely to be. DP snapshot intervals can range from once every 5 minutes to once every 24 hours.

A protection group is a group of VMs that have a common DP snapshot schedule and quiescence parameter value

Setting up DP snapshots requires two existing clusters running HX Data Platform version 2.5 or later. Both clusters must be on the same HX Data Platform version. Use HyperFlex Connect to complete the setup.

First, each cluster is set up for local replication networking. Use HX Connect to provide a set of IP addresses to be used by local cluster nodes to replicate to the remote cluster. HX Connect creates VLANs through UCS Manager, for dedicated local replication network use.



Note When this option is chosen in HX Connect, UCSM is configured only when both UCS Manager and fabric interconnect are associated with the HyperFlex cluster. When UCSM and FI are not present, you must enter the VLAN ID, and not select UCSM configuration in HX Connect.

The two clusters, and their corresponding existing relevant datastores must be explicitly paired. The pairing setup can be completed using HX Connect from one of the two clusters. This requires administrative credentials of the other cluster.

Virtual machines can be protected (or have their existing protection attributes modified) by using HX Connect at the cluster where they are currently active.

HX Connect can monitor the status of both incoming and outgoing replication activity on a cluster.

After a disaster, a protected VM can be recovered and run on the cluster that serves as the DP snapshot recovery site for that VM.

Replication and Recovery Considerations

The following is a list of considerations when configuring virtual machine replication and performing disaster recovery of virtual machines.



Note Cisco HX Data Replication is not supported with HyperFlex clusters that have Hardware Acceleration Cards installed.

- **Administrator**—All replication and recovery tasks, except monitoring, can only be performed with administrator privileges on the local cluster. For tasks involving a remote cluster, both the local and remote user must have administrator privileges and should be configured with the vCenter SSO on their respective clusters.
- **Minimum and Recommended Bandwidth**—Beginning with HX Release 4.0(2a) the minimum bandwidth can be configured to be 10 Mb for smaller size deployments. The replication network link should also be reliable and have sustained minimum symmetric bandwidth same as configured in a HyperFlex DR network. This should not be shared with any other applications on an Uplink or Downlink.
- **Maximum Latency**—Maximum latency supported is 75ms between two clusters.

If you are scheduling to run multiple replication jobs at the same time, for example 32 as maximum supported by DR, and your bandwidth (50Mbs) is low and latency (75ms) high, it is possible that some jobs will error out until bandwidth becomes available. If this situation occurs, retry the job, increase bandwidth or reduce the concurrency by staggering the replications.

During this situation, unprotect operations can take longer than expected.

- **Network Loss**—When there is a packet loss in data transmission across two sites, protection and recovery operations will have unexpected results. The transmission should be reliable for these features to work as expected.
- **Storage Space**—Ensure that there is sufficient space on the remote cluster to support the DR replication schedule. The protected virtual machines are replicated (copied) to the remote cluster at every scheduled

interval. Though storage capacity methods are applied (deduplication and compression), each replicated virtual machine will consume some storage space.

Not having sufficient storage space on the remote cluster can cause the remote cluster to reach capacity usage maximums. If you see **Out of Space** errors, see [Handling Out of Space Errors](#). Pause all replication schedules until you have appropriately adjusted the space available on the cluster. Always ensure that your cluster capacity consumption is below the space utilization warning threshold.

- **Supported Clusters**—Replication is supported between the following HyperFlex clusters:
 - 1:1 replication between HX clusters running under fabric interconnects.
 - 1:1 replication between All Flash and Hybrid HX cluster running under fabric interconnects.
 - 1:1 replication between 3-Node and 4-Node HX Edge and another 3-Node and 4-Node HX Edge cluster.
 - 1:1 replication between All Flash 3-Node and 4-Node Edge and Hybrid 3-Node and 4-Node HX Edge clusters.
 - 1:1 replication between 3-Node and 4-Node HX Edge and an HX cluster running under fabric interconnects.
 - Starting with HX release 4.5(2a), 1:1 replication with 2-Node HX Edge cluster.
- **Rebooting Nodes**—Do not reboot any nodes in the HX Cluster during any restore, replication, or recovery operation.
- **Thin Provision**—Protected VMs are recovered with thin provisioned disks irrespective of how disks were specified in the originally protected VM.
- **Protection Group Limitations**
 - Do not add VMs with ISOs or floppies to protection groups.

Protected Virtual Machine Scalability

- HX Release 3.5(x) supports the maximum limit of 200 VMs.
- **Non-HX Datastores**—Periodical replication fails on a protected a VM with storage on a non-HX datastore. To avoid the failure, unprotect the VM or remove non-HX storage from the VM.

Do not move protected VMs from HX datastores to non-HX datastores. If a VM is moved to a non-HX datastore through storage vMotion, unprotect the VM.
- **Templates**—Templates are not supported with disaster recovery.
- **Protection and Recovery of Virtual Machines with Snapshots**

When replication is enabled:

- **VM with no Snapshots**— The entire content of the VM is replicated.
- **VM with VMware redolog snapshots**— The entire content including the snapshot data is replicated. When a VM with redolog snapshots is recovered, all previous snapshots are preserved.
- **VM with HX native snapshots**—Only the latest data is replicated, and the HX native snapshot data is not replicated. When the VM is recovered, previous snapshots are not preserved.

- **Data Protection snapshots** are stored on the same datastore as the protected VMs. Manual deletion of DP snapshots is not supported. Deleting snapshot directories will compromise HX data protection and disaster recovery.



Caution To avoid deleting DP snapshots manually, it is important to remember that VMware does not restrict operations on datastores by the Admin user. In any VMware environment, datastores are accessed by the Admin via vCenter browser or by logging into the ESX host. Because of this, the snapshot directory and contents are browseable and accessible to Admins.

Other points for consideration include:

- **Location of the VMware vCenter**—If you delete a VM from VMware vCenter that is located on a “Other DRO” datastore pair, a recovery plan for this datastore pair fails during recovery. To avoid this failure, first unprotect the VM using the following command on one of the controller VM


```
stcli dp vm delete --vmid <VM_ID>
```
- **Name of the VM**—If you rename a VM from the vCenter, HyperFlex recovers at the previous name folder but registers the VM with the new name on the recovery side. Following are some of the limitations to this situation:
 - VMware allows a VMDK located at any location to be attached to a VM. In such cases, Hyperflex recovers the VM inside the VM folder and not at a location mapped to the original location. Also, recovery can fail if the VMDK is explicitly referenced in the `virtualmachine name.vmx` file by its path. The data is recovered accurately but there could be problems with registering the VM to vCenter. Correct this error by updating the `virtualmachine name.vmx` file name with the new path.
 - If a VM is renamed and a VMDK is added subsequently, the new VMDK is created at `[sourceDs] newVm/newVm.vmdk`. Hyperflex recovers this VMDK with the earlier name. In such cases, recovery can fail if the VMDK is explicitly referenced in the `virtualmachine name.vmx` file by its path. The data is recovered accurately but there could be problems with registering the VM to the Virtual Center. Correct this error by updating the `virtualmachine name.vmx` file name with the new path.
- **Network Ports**—The comprehensive list of ports required for component communication for the HyperFlex solution is located in Appendix A of the [HX Data Platform Security Hardening Guide](#).

Replication Network and Pairing Considerations

A replication network must be established between clusters that are expected to use replication for DP snapshots. The Replication network is created to isolate inter-cluster replication traffic from other traffic within each cluster and site.

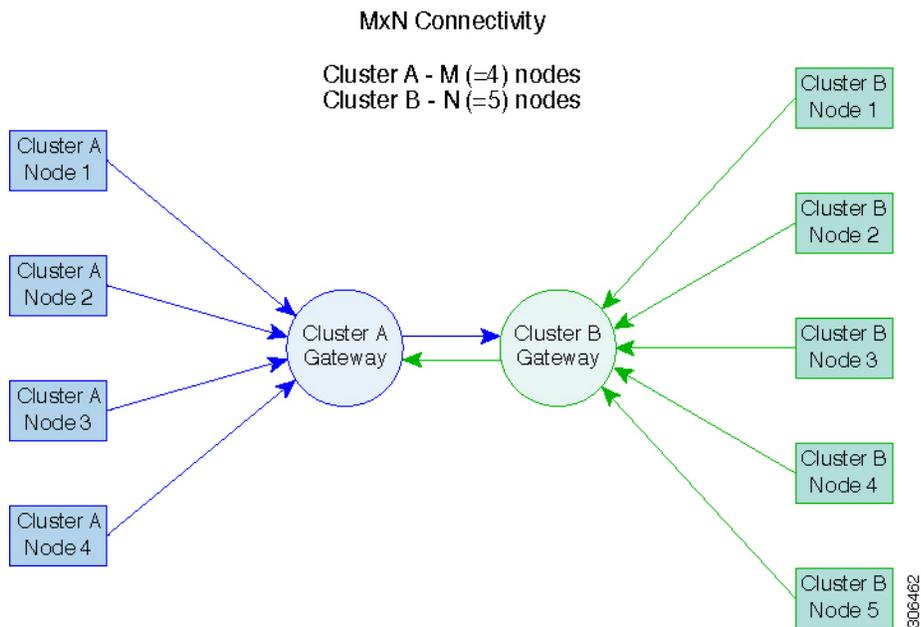
The following is a list of pre-validation checks necessary for pairing:

- Verify and ensure that peer credentials are working.
- Check the health of both clusters and proceed with pairing only when both the clusters are healthy.
- Verify and ensure that vCenter version is at least same or later than ESXi version at each cluster.

The following is a list of considerations when configuring replication network and pairing:

- To support efficient replication, all M nodes of cluster A have to communicate with all N nodes of cluster B, as illustrated in the *M x N connectivity between clusters* figure.
- To enable replication traffic between clusters to cross the site-boundary and traverse the internet, each node on Cluster A should be able to communicate with each node on Cluster B across the site boundary and the internet.
- The replication traffic must be isolated from other traffic within the cluster and the data center.
- To create this isolated replication network for inter-cluster traffic, complete these steps:
 - Create a replication network on each cluster.
 - Pair clusters to associate the clusters and establish M x N connectivity between the clusters.
- IP addresses, Subnet, VLAN, and Gateway are associated with the replication network of each cluster. Configure the corporate firewall and routers on both sites to allow communication between the clusters and the sites on TCP ports 9338,3049,9098,4049,4059.
- The comprehensive list of ports required for component communication for the HyperFlex solution is located in Appendix A of the [HX Data Platform Security Hardening Guide](#).

M*N Connectivity Between Clusters



Data Protection Terms

Interval—Part of the replication schedule configuration, used to enforce how often the protected VMs DP snapshot must be taken and copied to the target cluster.

Local cluster—The cluster you are currently logged into through HX Connect, in a VM replication cluster pair. From the local cluster, you can configure replication protection for locally resident VMs. The VMs are then replicated to the paired remote cluster.

Migration—A routine system maintenance and management task where a recent replication DP snapshot copy of the VM becomes the working VM. The replication pair of source and target cluster do not change.

Primary cluster—An alternative name for the source cluster in VM disaster recovery.

Protected virtual machine— A VM that has replication configured. The protected VMs reside in a datastore on the local cluster of a replication pair. Protected VMs have a replication schedule configured either individually or by inclusion in a protection group.

Protection group—A means to apply the same replication configuration to a group of VMs.

Recovery process—The manual process to recover protected VMs in the event the source cluster fails or a disaster occurs.

Recovery test—A maintenance task that ensures the recovery process will be successful in the event of a disaster.

Remote cluster—One of a VM replication cluster pair. The remote cluster receives the replication snapshots from the Protected VMs in the local cluster.

Replication pair—Two clusters that together provide a remote cluster location for storing the replicated DP snapshots of local cluster VMs.

Clusters in a replication pair can be both a remote and local cluster. Both clusters in a replication pair can have resident VMs. Each cluster is local to its resident VMs. Each cluster is remote to the VMs that reside on the paired local cluster.

DP snapshot—Part of the replication protection mechanism. A type of snapshot taken of a protected VM, which is replicated from the local cluster to the remote cluster.

Secondary cluster—An alternative name for the target cluster in VM disaster recovery.

Source cluster—One of a VM replication cluster pair. The source cluster is where the protected VMs reside.

Target cluster—One of a VM replication cluster pair. The target cluster receives the replicated DP snapshots from the VMs of the source cluster. The target cluster is used to recover the VMs in the event of a disaster on the source cluster.

Best Practices for Data Protection and Disaster Recovery

The requirement for an effective data protection and disaster recovery strategy based on the environment being protected cannot be overstated. The solution should be designed and deployed to meet or exceed the business requirements for both Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO) of the production VMs. The following are some of the points that must be considered when designing this strategy:

- The number of Service Level Agreements (SLA) necessary to comply with various categories of production workloads that may include mission critical, business critical, and important VMs.
- Detailed constructs of each SLA that may include RPO, RTO, the number or recovery points retained, requirements for off-site copies of data, and any requirements for storing backup copies on different media types. There may be additional requirements that include the ability to recover to a different environment such as a different location, different hypervisor or different private/public cloud.
- An ongoing testing strategy for each SLA which serves to prove that the solution meets the business requirements it was designed for.

Note that backups and backup copies must be stored external to the HyperFlex cluster being protected. For example, backups performed to protect VMs on a HyperFlex cluster should not be saved to a backup repository or a disk library that is hosted on the same HyperFlex cluster.

The built-in HyperFlex data protection capabilities are generalized into the following categories:

- **Data Replication Factor**—Refers to the number of redundant copies of data within a HyperFlex cluster. A data replication factor of 2 or 3 can be configured during data platform installation and cannot be changed. The data replication factor benefit is that it contributes to the number of cluster tolerated failures. See the section titled, [HX Data Platform Cluster Tolerated Failures](#) for additional information about the data replication factor.



Note Data Replication Factor alone may not fulfill requirements for recovery in the highly unlikely event of a cluster failure, or an extended site outage. Also, the data replication factor does not facilitate point-in-time recovery, retention of multiple recovery points, or creation of point-in-time copies of data external to the cluster.

- **HX Native Snapshots**—Operates on an individual VM basis and enables saving versions of a VM over time. A maximum of 31 total snapshots, including the HX SENTINEL snapshot, can be retained.



Note HX native snapshots alone may not fulfill requirements for recovery in the unlikely event of a cluster failure, or an extended site outage. Also, HX native snapshots do not facilitate the ability to create point-in-time copies of data external to the cluster. More importantly, unintentional deletion of a VM also deletes any HX native snapshots associated with the deleted VM.

- **Asynchronous Replication**—Also known as The HX Data Platform disaster recovery feature, it enables protection of virtual machines by replicating virtual machine DP snapshots between a pair of network connected HyperFlex clusters. Protected virtual machines running on one cluster replicate to the other cluster in the pair, and vice versa. The two paired clusters typically are located at a distance from each other, with each cluster serving as the disaster recovery site for virtual machines running on the other cluster.



Note Asynchronous Replication alone may not fulfill requirements for recovery when multiple point-in-time copies need to be retained on the remote cluster. Only the most recent snapshot replica for a given VM is retained on the remote cluster. Also, asynchronous replication does not facilitate the ability to create point-in-time copies of data external to either cluster.

It is recommended to first understand the unique business requirements of the environment and then deploy a comprehensive data protection and disaster recovery solution to meet or exceed those requirements.

Protecting Virtual Machines Overview

To protect a virtual machine (VM), specify the following protection attributes:

- Replication interval, at which DP snapshots are created for replication.
- A start time (within the next 24 hours), which specifies the first-time replication is attempted for that VM.
- Specify if the DP snapshot should be taken with the VM quiesced or not. Proper use of the quiesce option requires that VMware Tools are installed on the VM or VMs being protected.
- VMware Guest Tool for quiesce snapshot in Disaster Recovery is supported. Install the most recent VMware Guest Tool Service or verify that the existing service is up-to-date.



Note Using third-party guest tool (open-vm-tools) usage is allowed, but not supported by HX.

Protection attributes can be created and assigned to protection groups. To assign those protection attributes to VMs, they can be added to a protection group.

For example, there are three different SLAs: gold, silver, and bronze. Set up a protection group for each SLA, with replication intervals such as 5 or 15 minutes for gold, 4 hours for silver, and 24 hours for bronze. Most VMs can be protected by simply adding them to one of the three already created protection groups.

To protect VMs, you can choose from the following methods:



Note When you select multiple VMs, you must add them to a protection group.

- **Independently**—Select one VM and configure protection. Set the replication schedule and the VMware quiesce option for the specific VM. Changes to the replication settings will only affect the independently protected VM. The VM is not included in a protection group.
- **Existing protection group**—Select one or more VMs and add them to an existing protection group. The schedule and VMware quiesce option settings are applied to all of the VMs in the protection group. If the protection group settings are changed, the changes are applied to all of the VMs in the protection group.
- **New protection group**—Select two or more VMs and choose to create a new protection group. Define the protection group name, schedule, and VMware quiesce option settings. These settings are applied to all the VMs in the protection group. If the protection group settings are changed, the changes are applied to all of the VMs in the protection group.

Data Protection Workflow

To protect VMs and their data using DP snapshots and replication, perform the following steps:

- Configure two clusters and pair them to each other, to support the replication network activity.

- Assign a replication schedule to the VMs to set the frequency (interval) for creating DP snapshots on the source cluster and replicate them to the target cluster. Replication schedules can be assigned to individual VMs and to protection groups.

Replication Workflow

1. Install HX Data Platform, create two clusters.
2. Create at least one datastore on each cluster.
3. Log in to HX Connect.
4. Before creating the replication network, verify the IP addresses, subnet mask, VLAN, gateway, and IP range to be used for the replication network. After the replication network is created, validate connectivity within the cluster over the replication network.
5. The default value of MTU is 1500. If the HyperFlex cluster uses OTV or other tunneling mechanisms, ensure choosing an MTU value which will work for inter-site or inter-cluster connectivity.
6. Configure cluster replication network on each cluster. The replication network information is unique to each cluster.

Specify the subnet, gateway, range of IP addresses, bandwidth limit for dedicated use by the replication network. HX Data Platform configures a VLAN through UCS Manager for both clusters.
7. An intra-cluster network test is performed to validate connectivity between the nodes in the cluster, after the replication network is configured. If the intra-cluster network test fails, the replication network configuration is rolled back. Reconfigure the replication network after resolving any issues.
8. Before creating the replication pair, ensure that you have updated the corporate network to support this pairing.
9. Create a replication pair from one cluster to the other, connecting the two clusters. After the replication pair is created, a test of the inter-cluster pair network is performed to validate bidirectional connectivity between the clusters. Set the datastore mapping(s) from both clusters.
10. Optionally, you can create protection groups.
 - Set the schedule. Each protection group must have one schedule.
 - Create multiple protection groups if you want to have various replication intervals (schedules) for different VMs. A VM can only belong to one protection group.
11. Select VMs to protect, as individual virtual machines or VMs assigned to protection groups.
12. Set protection, do the following:
 - a. Select one or more VMs. Click Protect.
 - b. From the Protect VM wizard, the options are:
 - Protect a single VM with an existing protection group.
 - Protect a single VM independently.

Set the schedule.
 - Protect multiple VMs with an existing protection group.

- Protect multiple VMs with a new protection group.
Create new protection group and set schedule.

Configuring the Replication Network in HX Connect

Before a replication pair can be configured, the replication network has to be configured on both the local and remote cluster. Complete the configuration on the local cluster, then log in to the remote cluster and complete the configuration there.

Before you begin

Ensure that the following prerequisites are met, before configuring the replication network:

- A minimum of $N + 1$ IP addresses is required, where N is the number of converged nodes. An IP subnet spanning these new IP addresses, the gateway, and VLAN associated with this subnet is also required.
- To accommodate future cluster expansion, ensure that there are sufficient IP addresses in the subnet provided, for later use. Any new nodes in the expanded cluster would also need to be assigned IP addresses for replication. The subnet provided in the previous step should span the potentially new IP range as well.
- Additional IP-pool ranges can be added to the network later, however IP-pools already configured in the replication network cannot be modified.
- Make sure that the IP addresses to be used for the replication network are not already in use by other systems.
- Before creating the replication network, verify IP addresses, Subnet, VLAN, and Gateway to be used for the replication network.

Step 1 Log in to HX Connect as a user with administrator privileges.

Step 2 Select **Replication > Replication Configuration > Configure Network**.

Note You can only configure the replication network once. Once configured, you can edit the available IP addresses and the networking bandwidth.

Step 3 In the **Configure Replication Network** dialog box, under the **Configure Replication Network VLAN Configuration** tab, enter the network information.

UI Element	Essential Information
Select an existing VLAN radio button	Click this radio button to add an existing VLAN. If you manually configured a VLAN for use by the replication network through Cisco UCS Manager, enter that VLAN ID.
Create a new VLAN radio button	Click this radio button to create a new VLAN.

UI Element	Essential Information
VLAN ID field	<p>Click the up or down arrows to select a number for the VLAN ID or type a number in the field.</p> <p>This is separate from the HX Data Platform Management traffic network and Data traffic network.</p> <p>Important Be sure to use a different VLAN ID number for each HX Storage Cluster in the replication pair.</p> <p>Replication is between two HX Storage Clusters. Each HX Storage Cluster requires a VLAN dedicated to the replication network.</p> <p>For example, 3.</p> <p>When a value is added, the default VLAN Name is updated to include the additional identifier. The VLAN ID value does not affect a manually entered VLAN name.</p>
VLAN Name field	<p>This field is automatically populated with a default VLAN name when the Create a new VLAN radio button is selected. The VLAN ID is concatenated to the name.</p>
<p>For Stretched Cluster, provide Cisco UCS Manager credentials for primary and secondary FIs (site A and site B). For normal cluster, provide Cisco UCS Manager credential for single FI.</p>	
UCS Manager host IP or FQDN field	<p>Enter Cisco UCS Manager FQDN or IP address.</p> <p>For example, <i>10.193.211.120</i>.</p>
Username field	<p>Enter administrative username for Cisco UCS Manager.</p>
Password field	<p>Enter administrative password for Cisco UCS Manager.</p>

Step 4 Click **Next**.

Step 5 In the **IP & Bandwidth Configuration** tab, set the network parameters and the replication bandwidth.

UI Element	Essential Information
Subnet field	<p>Enter the subnet for use by the replication network in network prefix notation. The subnet is separate from the HX Data Platform Management traffic network and Data traffic network.</p> <p>Format example: <code>x.x.x.x/<number of bits></code> <code>1.1.1.1/20</code></p>
Gateway field	<p>Enter the gateway IP address for use by the replication network. The gateway is separate from the HX Data Platform Management traffic network and Data traffic network.</p> <p>For example, <i>1.2.3.4</i>.</p> <p>Note The gateway IP address must be accessible even if the disaster recovery is being setup for a flat network.</p>

UI Element	Essential Information
IP Range field	<p>Enter a range of IP addresses for use by the replication network.</p> <ul style="list-style-type: none"> The minimum number of IP addresses required is the number of nodes in your HX Storage Cluster plus one more. For example, if you have a 4 node HX Storage Cluster, enter a range of at least 5 IP addresses. The from value must be lower than the to value. For example, <i>From 10.10.10.20 To 10.10.10.30</i>. If you plan to add nodes to your cluster, include sufficient number of IP addresses to cover any additional nodes. You can add IP addresses at any time. <p>Note The IP address range excludes compute-only nodes.</p>
Add IP Range button	<p>Click to add the range of IP addresses entered in IP Range <small>From</small> and <small>To</small> fields.</p>

UI Element	Essential Information
<p>Set Replication Bandwidth Limit check box</p>	<p>Enter the maximum network bandwidth that the replication network is allowed to consume for inbound and outbound traffic. Acceptable value is 33,000 Mbits/sec.</p> <p>The default value is <code>unlimited</code>, which sets the maximum network bandwidth to the total available to the network.</p> <p>The replication bandwidth is used to copy replication snapshots from this local HX Storage Cluster to the paired remote HX Storage Cluster.</p> <p>Note</p> <ul style="list-style-type: none"> • At lower bandwidth (typically, lesser than 50 Mbps), the replications of multiple VMs may exit without executing the replication process due to high data transfer rate. To overcome this issue, either increase the bandwidth or stagger VM replication schedule so that VMs do not replicate in the same window. • The bandwidth setting must be close to the link speed. The bandwidth setting for the clusters in the pair must be same. • The set bandwidth is applicable only for the incoming and outgoing traffic of the cluster to which the bandwidth is set to. For example, setting the bandwidth limit as 100Mb means that the 100Mb is set for incoming traffic and 100Mb is set for outgoing traffic. • The set bandwidth limit must not exceed the physical bandwidth. • The bandwidth configured must be same on both sites of the disaster recovery environment. • The allowed low bandwidth is 10Mb and the maximum latency supported with 10Mb is 75ms. If the initial replication of VMs fails due to lossy network or unstable HX clusters, the VM replication will be initiated again in the next schedule as a fresh replication job. In this case, you have to size the schedule accordingly to protect VMs.

UI Element	Essential Information
Set non-default MTU check box	<p>Default MTU value is 1500.</p> <p>Select the check box to set a custom MTU size for the replication network. MTU can be set in the range 1024 to 1500.</p> <p>Note</p> <ul style="list-style-type: none"> • Ensure to use the same MTU value on both sides of the cluster. • After configuring the cluster, if the MTU value needs to be changed, the existing replication network configuration will need to be removed. The replication network can then be configured with the correct MTU value. <p>To reconfigure the replication network with a new MTU value, do the following:</p> <ol style="list-style-type: none"> a. Delete the replication pair if it is already configured. b. In the Web CLI of HX Connect, execute the following command: <pre>stcli drnetwork cleanup</pre> c. After completion of the <code>drnetwork cleanup</code> task, configure the replication network.

Note When you use an existing VLAN for replication network, the replication network configuration fails. You must add the self-created replication VLAN to the management vNIC templates in Cisco UCS Manager.

Step 6 Click **Next**.

Step 7 In the **Test Configuration** tab, check the replication network configuration.

Step 8 Click **Configure**.

What to do next

- Be sure to configure the replication network on both HX Storage Clusters for the replication pair.
- After the replication network is created on the cluster, each converged node on the cluster would be configured with an IP address on the eth2 interface.
- Check for duplicate IP assignment using `'arping'`.

For example: `arping -D -b -c2 -I ethX $replicationIP` (replace ethX and $replicationIP with actual values).`

If there is a duplicate IP assignment, it is necessary to remove the replication network assignments.

To reconfigure the replication network with proper IP assignment, do the following:

1. Delete the replication pair if it is already configured.
2. In the Web CLI of HX Connect, execute the following command:

```
stcli drnetwork cleanup
```

3. After completion of the `drnetwork cleanup` task, configure the IP address in the replication network.

Test Local Replication Network

To perform an intra-cluster replication network test, do the following:

-
- Step 1** Log in to HX Connect.
- Enter the HX Storage Cluster management IP address in a browser. Navigate to `https://<storage-cluster-management-ip>`.
 - Enter the administrative username and password.
 - Click **Login**.
- Step 2** In the Navigation pane, click **Replication**.
- Step 3** From the **Actions** drop-down list, select **Test Local Replication Network**.
- Step 4** Click **Run Test**.
- Step 5** On the **Activity** page, you can view the progress of the *Test Replication Network* task.
-

Editing the Replication Network

When you expand a HX Cluster that has replication configured, ensure that you have sufficient IP addresses available for the replication network. The replication network requires dedicated IP addresses, one for every node in the cluster plus one more. For example, in a 3 node cluster, four IP addresses are required. If you are adding one more nodes to the cluster, a minimum of five IP addresses are required.

To edit the replication network to add IP addresses, perform the following tasks:

-
- Step 1** Log in to HX Connect as administrator.
- Step 2** In the Navigation pane, Select **Replication**.
- Step 3** From the **Actions** drop-down list, select **Edit Replication Network**.
- Step 4** In the **Edit Network Configuration** dialog box, you can edit the range of IPs to use and set the replication bandwidth limit for replication traffic. The replication network subnet and gateway are displayed for reference only and cannot be edited.

UI Element	Essential Information
Subnet field	The subnet that is configured for the replication network in network prefix notation. This value cannot be edited.
Gateway field	The gateway that is configured for the replication network. This is value cannot be edited.

UI Element	Essential Information
IP Range field	<p>Enter a range of IP addresses for use by the replication network.</p> <ul style="list-style-type: none"> The minimum number of IP addresses required is the number of nodes in the HX Storage Cluster plus one more. For example, if the HX Storage Cluster has 4 nodes, the IP Range must be at least 5 IP addresses. The from value must be lower than the to value. For example, <i>From 10.10.10.20 To 10.10.10.30</i>. You can add IP addresses at any time. If you plan to add nodes to your cluster, include sufficient number of IP addresses to accommodate any additional nodes. <p>Note The IP address range excludes compute-only nodes.</p>
Add IP Range field	<p>Click to add the range of IP addresses that are entered in IP Range From and To fields.</p>
Set replication bandwidth limit check box (Optional)	<p>Enter the maximum network bandwidth that the replication network is allowed to consume for outbound traffic.</p> <p>Valid Range: 10 to 10,000. The default is <code>unlimited</code>, which sets the maximum network bandwidth to the total available replication network bandwidth.</p> <p>The replication bandwidth is used to copy DP snapshots from the local HX Storage Cluster to the paired remote HX Storage Cluster.</p>

Step 5 Click **Save Changes**.

The replication network is now updated. Any additional replication network IP addresses are available for new nodes should they be added to the storage cluster. Replication traffic adjusts to any changes made to the bandwidth limit.

Replication Pair Overview

Creating a replication cluster pair is a prerequisite for configuring VMs for replication. After two (2) HX clusters are paired, map the datastore on the remote cluster to a datastore on the local cluster.

Mapping datastore A on HX cluster 1 with a datastore B on HX cluster 2 enables any VM on HX cluster 1 that resides in datastore A and is configured for replication to be replicated to datastore B on HX cluster 2. Similarly, any VM on cluster 2 that resides in datastore B and is configured for replication to be replicated to datastore A on HX cluster 1.

Pairing is strictly 1-to-1. A cluster can be paired with no more than one other cluster.

Mapping is a strict 1-to-1 relationship. A datastore on a paired HX cluster can be mapped to no more than one datastore on the other HX cluster. Note that there can be multiple mapped datastores. For example, datastore A on HX cluster 1 mapped to datastore B on HX cluster 2, and datastore C on HX cluster 1 mapped to datastore D on HX cluster 2.

Creating a Replication Pair

The replication pair defines the two halves of the protection network. The HX Storage Cluster you are logged into is the local cluster, the first half of the pair. Through this dialog, you identify another HX Storage Cluster, the second half of the pair, the remote cluster. To ensure the storage component, map the replication pair to datastores on each HX Storage Cluster. After the replication pair is configured, you can begin protecting virtual machines. See the **Virtual Machines** tab.



Important

When pairing or mapping clusters at different versions, cluster pairing must be initiated on the 3.5 cluster and then the datastore mapping must be initiated from the 3.0 cluster.

- Cluster pairing must be initiated only from a 3.5 cluster to a 3.0 cluster.
- Datastore mapping must be initiated only from a 3.0 cluster to a 3.5 cluster.

Before you begin

- Create a datastore on both the local and the remote cluster.
- Configure the replication network.

Step 1 From HX Connect, log in to either the local or remote cluster as a user with administrator privileges. Select **Replication > Replication Pairs > Create Replication Pair**.

Step 2 Enter a **Name** for the replication pair and click **Next**.

Enter a name for the replication pairing between two HX Storage Clusters. This name is set for both the local and remote cluster. The name cannot be changed.

Step 3 Enter the **Remote Connection** identification and click **Pair**.

Once the Test Cluster Pair job is successful, you can proceed to the next step. On the Activity page, you can view the progress of the Test Cluster Pair job.

UI Element	Essential Information
Management IP or FQDN field	Enter the IP address or fully qualified domain name (FQDN) for the management network on the remote HX Storage Cluster. For example: <i>10.10.10.10</i> .
User Name and Password fields	Enter vCenter single sign-on or cluster specific administrator credentials for the remote HX Storage Cluster.

HX Data Platform verifies the remote HX Storage Cluster and assigns the replication pair name.

Note Virtual machines to be protected must reside on one of the datastores in the replication pair.

Step 4 Set the **Datastore Mapping** from both clusters and click **Next**.

- Note**
- The virtual machines to be protected must on the datastores you select. Moving virtual machines from the configured datastores for the replication pair, also removes protection from the virtual machines.
 - Moving virtual machine to another paired datastore is supported. If the VMs are moved to unpaired datastore, replication schedule fails.

UI Element	Essential Information
Local Datastore column	List of the configured datastores on this cluster, the local HX Storage Cluster. Map one local datastore to one remote datastore.
Remote Datastore column	Pair the datastores between the HX Storage Clusters. From the desired Local Datastore row, select a datastore from the Remote Datastore pull-down menu. This selects both the remote and local datastores in a single action.

Step 5 Review the Summary information and click **Map Datastores**.

UI Element	Essential Information
Datastore field	The selected datastore on this local HX Storage Cluster.
Target Datastore field	The datastore on the remote HX Storage Cluster where the replication snapshot is copied to.

Test Remote Replication Network

To test the pairing between clusters in a remote replication network, do the following:

Step 1 Log in to HX Connect.

- a) Enter the HX Storage Cluster management IP address in a browser. Navigate to `https://<storage-cluster-management-ip>`.
- b) Enter the administrative username and password.
- c) Click **Login**.

Step 2 In the Navigation pane, click **Replication**.

Step 3 From the **Actions** drop-down list, select **Test Remote Replication Network**.

Field	Description
MTU Test Value	Default MTU value is 1500. MTU can be set in the range 1024 to 1500. Note <ul style="list-style-type: none"> • After configuring the cluster, if the MTU value needs to be changed, the existing replication network configuration will need to be removed. The replication network can then be configured with the correct MTU value. Be sure to use the same MTU value on both sides of the cluster.

Step 4 Click **Run Test**.

Step 5 On the **Activity** page, you can view the progress of the *Replication Pair Network Check* task.

Editing a Replication Pair

Editing a replication pair is changing the datastores for the replication pair.

- Step 1** Login to HX Connect as an administrator.
- Step 2** Select **Replication > Replication Pairs > Edit**.
- Step 3** Select the local or remote datastore and click **Finish**.

UI Element	Essential Information
Local Datastore column	List of the configured datastores on this cluster, the local HX clusters. Map one local datastore to one remote datastore.
Remote Datastore column	Pair the datastores between the HX clusters. <ol style="list-style-type: none"> To change the local datastore selection, remove the mapping to the current local datastore. From the pull-down menu in the Remote Datastore column, select Do not map this datastore. From the desired Local Datastore row, select a datastore from the Remote Datastore pull-down menu. This selects both the remote and local datastores in a single action.

Deleting a Replication Pair

Delete a replication pair on the local and remote clusters.

Select **Replication > Replication Pairs > Delete**.

Before you begin

On both the local and remote HX clusters, remove dependencies from the replication pair.

Log in to the local and the remote HX storage cluster and perform the following:

- Unprotect all virtual machines. Remove virtual machines from protection groups.
- Remove protection groups. If the protection group does not have a VM, deleting protection group is not required.

- Step 1** Log in to HX Connect as an administrator.
- Step 2** Unmap the datastores in the replication pair.
- Select **Replication > Replication Pairs > Edit**.
After the test cluster pair job is successful, you can proceed to the next step. You can view the progress of the Test Cluster Pair job on the Activity page.
 - From the **Edit Replication Pair** dialog box, select **Do not map this datastore** from the **Remote Datastore** menu.

UI Element	Essential Information
Local Datastore column	List of the configured datastores on this cluster, the local HX clusters. Map one local datastore to one remote datastore.
Remote Datastore column	Pair the datastores between the HX clusters. <ol style="list-style-type: none"> To change the local datastore selection, remove the mapping to the current local datastore. From the pull-down menu in the Remote Datastore column, select Do not map this datastore. From the desired Local Datastore row, select a datastore from the Remote Datastore pull-down menu. This selects both the remote and local datastores in a single action.

- c) Ensure all the possible remote datastores are set to **Do not map this datastore**.
d) Click **Finish**.

Step 3 Select **Replication > Replication Pairs > Delete**.

Step 4 Enter administrator credentials for the remote cluster and click **Delete**.

UI Element	Essential Information
User Name field	Enter the administrator user name for the remote HX Storage Cluster.
Password field	Enter the administrator password for the remote HX Storage Cluster.

Creating a Protection Group

A protection group is a group of VMs with the same replication schedule and VMware Tools quiescence settings.

Protection groups are created on the HX cluster that the administrative user is logged on to. Protection groups provide protection to the VMs that are members of a given protection group. If protection groups have protected virtual machines that replicate to the remote cluster, they are listed in HX Connect.



Note The administration of protection groups can only be performed from the local cluster where it was created.

Before you begin

- Ensure that replication network and replication pair are configured.
- Install the most recent VMware Guest Tool Service or verify that the existing service is up-to-date.

Step 1 Log in to HX Connect as an administrator.

Step 2 Select **Replication > Protection Groups > Create Protection Group**.

Step 3 Enter the information in the dialog fields.

UI Element	Essential Information
Protection Group Name field	Enter a name for the new protection group for this HX cluster. Protection groups are unique to each HX cluster. The name is referenced on the remote HX cluster, but not editable on the remote HX cluster. Multiple protection groups can be created on each HX cluster.
Protect virtual machines in this group every field	Select how often the virtual machines are to be replicated to the paired cluster. The pull-down menu options are: 5 minutes, 15 minutes, 30 minutes, 1 hour, 90 minutes, 2 hours, 4 hours, 8 hours, 12 hours, and 24 hours. The default value is 1 hour.
Start protecting the virtual machines immediately radio button	Select this radio button if you want the first replication to start immediately after you add the first virtual machine to this protection group.
Start protecting the virtual machines at radio button	Select this radio button if you want to set a specific time for the first replication operation to start. Before you start replication ensure: <ul style="list-style-type: none"> • At least one virtual machine is added to the protection group. • The scheduled start time is reached. To specify the protection operation start time: <ol style="list-style-type: none"> a. Check the Start protecting the virtual machines at radio button. b. Click in the time field and select an hour and minute. Then click out of the field. <p>Cluster time zone and Current time on cluster are references to help you to choose the appropriate replication start time. Start time is based on the local cluster clock. For example:</p> <p>10 hours, 3 minutes from now with Current time on cluster, 1:56:15PM, means that the first replication occurs at 11:59:00PM.</p> <p>The <i>hours, minutes from now</i> states when the first replication will occur. This is updated when you change the time field setting.</p>
Use VMware Tools to quiesce the virtual machine check box	Select this check box to take quiesced DP snapshots. Leaving the check box in an unchecked state will take crash consistent DP snapshots. This only applies to virtual machines with VMware Tools installed.

Step 4 Click **Create Protection Group**.

HX Data Platform adds the new group to the **Protection Groups** tab. This protection group is available to protect virtual machines on this cluster.

Step 5 Click the **Replication > Protection Groups** to view or edit the new protection group.

If the number of VMs is zero (0), add virtual machines to the new protection group to apply the replication schedule configured in the protection group.

Editing Protection Groups

Change the replication interval (schedule) for the virtual machines in the protection group. To edit the protection groups, perform the following steps:

- Step 1** Login to HX Connect as an administrator.
- Step 2** Select **Replication > Protection Groups > Edit Schedule**.
- Step 3** Edit the information in the dialog fields.

UI Element	Essential Information
Protect virtual machines in this group every field	Use the pull-down list to select how often the virtual machines are to be replicated to the paired cluster. List values are: 5 minutes, 15 minutes, 30 minutes, 1 hour, 90 minutes, 2 hours, 4 hours, 8 hours, 12 hours, and 24 hours.
Use VMware Tools to quiesce the virtual machine check box	Select the check box to take quiesced DP snapshots. The checkbox is unchecked by default; leaving the check box unchecked, takes crash consistent DP snapshots. This only applies to virtual machines with VMware Tools installed.

- Step 4** Click **Save Changes** to save the interval and VMware Tools quiescence settings for the protection group. View the Protection Groups tab to verify the interval frequency.

Deleting Protection Groups

Before you begin

Remove all virtual machines from the protection group.

- Step 1** Select **Replication > Protection Groups > *protection_group_name***
- Step 2** Click **Delete**. Click **Delete** in the verification pop-up.

Protecting Virtual Machines with an Existing Protection Group

This task describes how to protect multiple virtual machines using an existing protection group.

Using an **Existing protection group**—Select one or more virtual machines and add them to an existing protection group. The schedule and VMware quiesce option settings are applied to all the virtual machines in

the protection group. When the protection group settings are changed, the changes are applied to all the virtual machines in the protection group.

Before you begin

Replication network and replication pair configured.

Create protection group prior to adding the virtual machines.

Step 1 Log in to HX Connect with administrator privileges and select **Virtual Machines**.

This lists the virtual machines on the local HX cluster.

Step 2 Select one (1) or more unprotected VMs from the list.

Click in the virtual machine row to select it. As you click a virtual machine row, the corresponding virtual machine check box is selected.

Step 3 Click **Protect**.

The **Protect Virtual Machines** wizard, **Protection Group** page is displayed.

Step 4 Click the radio button, **Add to an existing protection group**

UI Element	Essential Information
Set the protection parameters table	Verify the selected virtual machine Name . Use the Storage Provisioned and Storage Used to check for sufficient storage resource availability on the remote HX cluster.
Add to an existing protection group radio button	Select an existing protection group from the pull-down list. The interval and schedule settings of the protection group are applied to the selected VM or VMs.
Create a new protection group radio button	Enter a name for the new protection group for this local cluster. Protection groups are unique to each cluster. The name is referenced on the remote cluster, but not editable on the remote cluster. You can create multiple protection groups on each cluster.

Step 5 Select a protection group from the pull-down list and click **Next**.

Be sure the protection group you choose has the schedule interval desired.

The **Protect Virtual Machines** wizard, **Summary** page is displayed.

Step 6 Confirm the information in the **Summary** page and click **Add to Protection Group**.

The selected VM or VMs are added to the protection group. View the **Replication** or **Virtual Machines** pages to confirm that the VM or VMs have been added to the protection group.

Protecting Virtual Machines with a New Protection Group

This task describes how to protect multiple virtual machines by creating a new protection group.

Using a **New protection group**—Select VMs and choose to create a new protection group. Define the protection group name, schedule, start time, and VMware quiesce option settings. These settings are applied to all the virtual machines in the protection group. When the protection group settings are changed, the changes are applied to all the virtual machines in the protection group.

Before you begin

Replication network and replication pair configured.

Step 1 Login to HX Connect with administrator privileges and select **Virtual Machines**.

This lists the virtual machines on the local HX cluster.

Step 2 Select one or more unprotected VMs from the list.

Click in the virtual machine row to select it. As you click a virtual machine row, the corresponding virtual machine checkbox is selected.

Step 3 Click **Protect**.

The **Protect Virtual Machines** wizard, **Protection Group** page is displayed.

Step 4 Click the radio button, **Create a new protection group**, add a name for the protection group, and click **Next**.

The **Protection Schedule Wizard Page** wizard page is displayed.

Step 5 Complete the schedule and VMware quiesce option, as needed, and click **Next**.

UI Element	Essential Information
Protect virtual machines in this group every field	Select how often the virtual machines are to be replicated to the paired cluster. The default value is every 1 hour.
Start protecting the virtual machines immediately radio button	Select this radio button if you want the first replication to start immediately after you add the first virtual machine to this protection group.

UI Element	Essential Information
<p>Start protecting the virtual machines at radio button</p>	<p>Select this radio button if you want to set a specific time for the first replication to start. To start replication requires:</p> <ul style="list-style-type: none"> • At least one virtual machine is added to the protection group. • The scheduled start time is reached. <p>To specify the protection operation start time:</p> <ol style="list-style-type: none"> a. Check the Start protecting the virtual machines at radio button. b. Click in the time field and select an hour and minute. Then click out of the field. <p>The <i>hours, minutes from now</i> states when the first replication will occur. This is updated when you change the time field setting.</p> <p>Cluster time zone and Current time on cluster are references to help you to choose the appropriate replication start time. Start time is based on the local cluster clock. For example:</p> <p>10 hours, 3 minutes from now with Current time on cluster, 1:56:15PM, means that the first replication occurs at 11:59:00PM.</p>
<p>Use VMware Tools to quiesce the virtual machine check box</p>	<p>Click the check box to take quiesced DP snapshots. An unchecked check box takes crash consistent DP snapshots. The check box is unchecked by default.</p> <p>This only applies to virtual machines with VMware Tools installed.</p>

The **Protect Virtual Machines** wizard, **Summary** page is displayed.

Step 6 Confirm the information in the **Summary** page and click **Add to Protection Group**.

Review the summary content to confirm the settings to apply to the selected virtual machines.

- Name of the protection group
- Number of virtual machines to protect
- Names of virtual machines
- Storage provisioned for each virtual machine
- Storage used (consumed) by each virtual machine

The selected VM or VMs are added to the protection group. View the **Replication** or **Virtual Machines** pages to verify that the VM(s) have been added to the protection group.

Protecting Individual Virtual Machines

This task describes how to protect a virtual machine (VM).

- **Independently**—Select one (1) VM and configure protection. Set the replication schedule and the VMware Tools quiesce option for the specific VM.

Changes to the replication settings only affect the independently protected VM. The VM is not a member of a protection group.

- **Existing protection group**—Select one or more VMs and add them to an existing protection group. The schedule and VMware Tools quiesce option settings are applied to all the VMs in the protection group. When the protection group settings are changed, the changes are applied to all VMs in the protection group.

Before you begin

Configure replication network and replication pair.

Step 1 Log in to HX Connect with administrator privileges and select **Virtual Machines**.

Step 2 Select one unprotected virtual machine from the list. Click in the virtual machine row to select it.

Click in the virtual machine row to select it. As you click a virtual machine row, the corresponding virtual machine check box is selected.

Step 3 Click **Protect**.

The **Protect Virtual Machine** dialog box is displayed.

Step 4 Complete the fields as needed.

UI Element	Essential Information
Add to an existing protection group radio button	Select an existing protection group from the pull-down list. The interval and schedule settings of the protection group are applied to this virtual machine. No additional configuration is required, click Protect Virtual Machine .
Protect this virtual machine independently radio button	Enables the interval, schedule options, and VMware Tools quiescence option for defining protection for this VM.
Protect this virtual machine every field	Select from the pull-down list how often the virtual machines are to be replicated to the paired cluster. The list values are: 5 minutes, 15 minutes, 30 minutes, 1 hour, 90 minutes, 2 hours, 4 hours, 8 hours, 12 hours, and 24 hours.
Start protecting the virtual machines immediately radio button	Select this radio button if you want the first replication to start immediately after you add the first virtual machine to this protection group.

UI Element	Essential Information
<p>Start protecting the virtual machines at radio button</p>	<p>Select this radio button if you want to set a specific time for the first replication to start. To start replication requires:</p> <ul style="list-style-type: none"> • At least one VM is added to the protection group. • The scheduled start time is reached. <p>To specify the protection operation start time:</p> <ol style="list-style-type: none"> a. Check the Start protecting the virtual machines at radio button. b. Click in the time field and select an hour and minute. Then click out of the field. <p>The <i>hours, minutes from now</i> states when the first replication will occur. This is updated when you change the time field setting.</p> <p>Cluster time zone and Current time on cluster are references to help you to choose the appropriate replication start time. Start time is based on the local cluster clock. For example:</p> <p>10 hours, 3 minutes from now with Current time on cluster, 1:56:15PM, means that the first replication occurs at 11:59:00PM.</p>
<p>VMware Tools to quiesce the virtual machine check box</p>	<p>To take quiesced DP snapshots, check the check box. The unchecked check box takes crash consistent DP snapshots. The check box is unchecked by default.</p> <p>This only applies to virtual machines with VMware Tools installed.</p>

Step 5 Click **Protect Virtual Machine**.

The VM status is updated in the **Virtual Machine** and the **Replication** page. Notice on the Replication page, that no Protection Group is listed for any VMs protected as individual VMs.

Replication is now enabled for this VM.

Unprotecting Virtual Machines



Note There is no need to unprotect VMs to pause replication for HX cluster activities. See [Pausing Replication, on page 37](#).



Note The Unprotect process may take several minutes to complete if the replication bandwidth is at 50 Mbps or lower and/or high latency (75ms and higher).

Step 1 Log in to HX Connect as an administrator.

- Step 2** Select **Virtual Machines**.
- Step 3** Select a protected virtual machine from the list. Click in the virtual machine row.
VMs can be unprotected one VM at a time.
- Step 4** Click **Unprotect** and click to confirm.
The state changes for the virtual machine from **protected** to **unprotected**.
-

Disaster Recovery Overview

Disaster recovery is performed when the source site is not reachable and you want to failover the VMs and the protected groups to the target cluster. The process of recovery recovers the VM on the target cluster. Recovering virtual machines is restoring a most recent replication snapshot from the recovery (target) cluster.

Software encryption must be enabled on clusters in both paired datastores to be able to protect VMs on encrypted datastores.

Testing VM recovery—The ability to test recovery without breaking replication. It can bring up your VM workload on the target to verify the contents of the VM.

Recovering virtual machines—Restoring a most recent replication snapshot from the target (recovery) cluster. Once you start Recovery, all scheduled replication will be stopped.

Planned migration—Performing planned migration pauses the replication schedule, creates and replicates a DP snapshot, and recovers on the target. Ownership is switched from the source to the target, and resumes replication on the target that is now the new source

Disaster Recovery and Reprotect—Recovers the VM on the target, switches the ownership from the source to the target, and resumes replication on the target that is now the new source.

Protecting VMs after disaster—In the event of a disaster, you may lose the source site altogether. After the recovery is performed protect the recovered VMs to a newer cluster.

Compatibility Matrix for Disaster Recovery Operations

The following compatibility matrix lists the DR operations that are supported when a cluster at HX Data Platform version 3.5(x) is paired with a cluster at HX Data Platform version 3.5(x) or 3.0(1x).

Feature	3.5(x) Paired With 3.5(x)	3.5(x) Paired With 3.0(1x)
Replication	√	√
Cluster Pairing	√	√
Datastore Mapping	√	√
Protection	√	√
Planned Migration (Single click using HX Connect)	√	X

Feature	3.5(x) Paired With 3.5(x)	3.5(x) Paired With 3.0(1x)
Planned Migration (Multi-step stcli or WebCLI and HX Connect for Recovery)	√	√
Test Recover using HX Connect	√	√
Recover using HX Connect	√	√
Re-protect using HX Connect	√	X
Re-protect (Multi-step stcli or WebCLI)	√	√

Testing Virtual Machine Recovery

Testing VM recovery gives you the ability to test recovery without breaking replication. It can bring up your VM workload on the target to verify the contents of the VM.



Note

- Testing recovery does not disrupt the running clusters. The intent is to verify, in the event of an actual disaster, that the VMs are recoverable.
- Using the HX Connect user interface, to test VM recovery, you can run a maximum of 10 tasks in a sequence without waiting for the previously submitted task to complete.

Before you begin

Before you begin the test VM recovery process, ensure the following:

- The target cluster is up and in good health.
- The protected VMs completed a recent replication to the target cluster. These replicated VMs are stored as DP snapshots on the target clusters.



Important

Only one copy of the test recovered VM can be made at any point. If you need to have another test recovered VM, please delete the previously created test recovered VM.

To the test VM recovery process perform the following steps:

Step 1 Log in to HX Connect on the target cluster as administrator.

Step 2 Navigate to **Replication > Remote VMs Tab > protected_vm**.

Step 3 To test the recovery process, click the **Test Recovery** button.

Note When configuring recovery settings, the following fields are auto-populated.

UI Element	Essential Information
Resource Pool drop-down list	Select a location for the test VM to be stored.
Folders drop-down list	Select a location for the test VM to be stored, for example: <ul style="list-style-type: none"> • Discovered Virtual Machine • HX Test Recovery
Power On/Off radio button	By default, the Power ON option is selected. The recovered VM is powered on or left off after it is created as per the selected option.
VM Name field	Enter a new name for the created test VM.
Test Networks radio button	Select which HX Storage Cluster network to use for transferring the data from the replication snapshot. Network options for example: <ul style="list-style-type: none"> • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network
Map Networks radio button	Select to create a map between the source and the target cluster networks. <ul style="list-style-type: none"> • Source Network—Network name at the source side on which the VM is connected. • Target Network—Select target network from the drop-down list, where the VM has to be connected.

Step 4 Click **Recover VM**.

Step 5 For VMs that are part of a protection group, perform a test recovery on each VM in the group.

Step 6 Verify the contents of the recovered VM.

Recovering Virtual Machines

Recovering VMs is restoring a most recent replication snapshot from the target (recovery) cluster.



Attention

- You may configure the folder, network, or resource pool parameters to be used during recovery, test recovery and migrate operations. If the global recovery setting is not configured, you will need to explicitly map individual VMs at the time of recovery.
- Recover VM is not supported between different vSphere versions. If the Target is at a lower version vSphere environment and does not support the hardware version of a protected VM on the primary, VM test recovery and recovery may fail. Cisco recommends to test recover each protected VM to validate the support on the target site.

Upgrade the target environment to enable recovery of protected VMs.

- Cancelling a recovery process (rolling back) is not supported. Attempting to cancel a recovery process changes all VMs in an unprotected ‘ready to recovery’ state.
- When running recovery on VMs, you may specify explicit network mapping when recovering the VMs to avoid unintentional network connections to recovered VMs.

You can skip specifying network mapping in the following cases:

- If the source VMs use vSphere Standard Switches and if all ESXi hosts on the recovery side have standard switch networks with the same name.
- If the source VMs use vSphere Distributed Switch (vDS) port groups and if the recovery site has identically named vDS port groups.
- If you want to specify network mapping, ensure that both the name and the type of the VM network matches between the source and the target.
- When running recovery on virtual machines that are individually protected, or that are in different protection groups, the maximum number of concurrent recovery operations on a cluster is 20.

Before you begin

Ensure the following:

- The target cluster is up and in good health.
- The protected VMs completed a recent replication to the target cluster. Replicated VMs are stored as DP snapshots on the target clusters.

On the target cluster, perform the following to conduct disaster recovery.

- Step 1** Log in to HX Connect as administrator.
- Step 2** Select **Replication** > > **Remote VMs tab** >> *protected_vm* and click **Recover**.
- Step 3** To recover the VM and build a new VM on the local cluster, click the **Recover VM** button.

Note When you configure recovery settings, the following fields are auto-populated.

UI Element	Essential Information
Resource Pool drop-down list	Select a location for the new VM to be stored.

UI Element	Essential Information
Folders drop-down list	Select a location for the new VM to be stored.
Power On/Off radio button	By default Power ON option is selected. The recovered VM is powered on or left off after it is created as per the selected option.
Map Networks	<p>Select to create a map between the source and target cluster networks.</p> <ul style="list-style-type: none"> • Source Network—Network name at the source side on which the VM is connected. • Target Network—Select target network from the drop-down list, where the VM has to be connected. <p>Network options for example:</p> <ul style="list-style-type: none"> • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

Step 4 Click **Recover VM**.

Step 5 Wait for the recovery to complete. View the recovered VM in the target vCenter.

Recovering Virtual Machines in Protection Groups

Step 1 Select a *protected-vm* and click **Recover**.

All VMs will be moved from the protection group and the selected VMs will be recovered. Recovered VMs show protection status as *Recovered* and the remaining (protection group) VMs show protection status as *Recovering*. The protection group will go in *Recovered* state and is not reusable. You can delete it from the primary site.

The recovered VMs are displayed in the **Standalone Protected VMs** subpane.

Step 2 Recover the remaining virtual machines from the **Standalone Protected VMs** subpane, which were a part of the protection group. See [Recovering Virtual Machines, on page 30](#) for more details.

Planned Migration

Performing planned migration pauses the replication schedule, replicates the most recent copy, recovers on the target, switches the ownership from the source to the target, and resumes replication on the target that is now the new source.



Attention This process cannot be rolled back.

- Step 1** Log in to HX connect of the target cluster. The target cluster is where the replicated DP snapshots were copied to.
- Step 2** On the target cluster, select **Replication > Remote VMs Tab > protected_vm**.
- Step 3** Click **Migrate**.

Note All the fields that are listed here are optional.

UI Element	Essential Information
Resource Pool drop-down list	Select a location for the new VM to be stored.
Folders drop-down list	Select a location for the new VM to be stored.
Power On/Off radio button	By default Power ON option is selected. The recovered VM is powered on or left off after it is created as per the selected option.
Map Networks	<p>Select to create a map between the source and target cluster networks.</p> <ul style="list-style-type: none"> • Source Network—Network name at the source side on which the VM is connected. • Target Network—Select target network from the drop-down list, where the VM has to be connected. <p>Network options for example:</p> <ul style="list-style-type: none"> • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

- Step 4** Monitor the progress on the **Activity** page.

Low Bandwidth and Temporary Packet Loss - In the event VM migration operation fails with an error message that contains "THRIFT_EAGAIN (timed out)", retry the VM Migration. The timeout error is due to temporary network congestion caused by bandwidth saturation or underlying network packet loss.

Migrating Virtual Machines in Protection Groups

Using the HX Connect user interface, to migrate VMs, you can run a maximum of 4 tasks in a sequence without waiting for the previously submitted task to complete.

- Step 1** Select a *protected-vm* and click **Migrate**.

All the VMs are now moved out from the protection group and are displayed in the **Standalone Protected VMs** subpane. Only the selected VM is recovered.

Step 2 Migrate the remaining virtual machines from the **Standalone Protected VMs** subpane, which were a part of the protection group. See [Planned Migration, on page 32](#) for more details.

Disaster Recovery and Re-protect

Performing disaster recovery recovers the VM on the target, switches the ownership from the source to the target, and resumes replication on the target that is now the new source. Disaster recovery is typically done when disaster occurs, and when you want to reverse the direction of protection.



Attention

- This process cannot be rolled back.
- For a single vCenter deployment, the Disaster Recovery workflow performed entirely through the HX Connect UI is not supported. To perform Disaster Recovery and Re-protect:
 1. Log in to vSphere Web Client Navigator of the primary site and remove the VM from the primary site to unregister the VM.
Right-click on the virtual machine and select **All vCenter Actions > Remove from Inventory**.
 2. Log in to HX Connect of the secondary site. Select **Replication > Remote VMs Tab > protected_vm**. Click **Recover**.
 3. When the primary site comes back up, log in to HX Connect of the secondary site. Select **Replication > Remote VMs Tab > protected_vm**. Click **Re-protect**.
 4. After Re-protect has completed successfully, log in to vSphere Web Client of the secondary site and manually register the VM.
 - a. Log in to vSphere Web Client Navigator. Select **Configuration > Storage**.
 - b. Right-click on the appropriate datastore and click **Browse Datastore**.
Navigate to the *virtualmachine name.vmx* file, right-click on the file and click **Add to Inventory**. Follow the wizard to manually register the VM.
- Using the HX Connect user interface, you can run a maximum of 5 re-protect tasks in a sequence without waiting for the previously submitted task to complete.

Step 1 Log in to HX connect of the source and the target. The target cluster is where the replicated DP snapshots were copied to. The source cluster is the cluster where the VMs reside.

Step 2 Select a VM from the remote VM list. Execute the Recover VM workflow on this cluster workflow.

Note If both the target and source clusters are on the same vCenter, then unregister the VM on the source cluster. This ensures that vCenter no longer has a record of the VM and it stops managing the VM, but it retains the data for the VM.

Step 3 Select **Replication > > Remote VMs tab >> protected_vm** and click **Recover**.

Step 4 To recover on the target VM and build a new VM on the local cluster, click the **Recover VM** button.

Complete the following fields in the **Recover VM on this cluster** dialog box.

UI Element	Essential Information
Resource Pool drop-down list	Select a location for the new VM to be stored.
Folders drop-down list	Select a location for the new VM to be stored.
Power On/Off radio button	By default the Power ON option is selected. The recovered VM is powered on or left off after it is created as per the selected option.
Map Networks	<p>Select to create a map between the source and target cluster networks.</p> <ul style="list-style-type: none"> • Source Network—Network name at the source side on which the VM is connected. • Target Network—Select target network from the drop-down list, where the VM has to be connected. <p>Network options for example:</p> <ul style="list-style-type: none"> • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

Step 5 Click **Recover VM**.

Step 6 On the target cluster, select **Replication > Remote VMs Tab > protected_vm**.

Step 7 Click **Re-protect**.

- Attention**
- If both the target cluster and source cluster are on the same vCenter, manually register the VM on the source cluster.
 - When the Re-protect task fails and in the HX Connect UI the **Re-protect** tab is not available, execute *stcli reverseprotect* to complete the Re-protect operation.

Protection status of the VM shows as **Protected**.

Step 8 After the original primary comes up, to migrate back to the primary do the following:

- On the target cluster, select **Replication > Remote VMs Tab > protected_vm**.
- Click **Migrate** to unregister the target VM and transfer the VM ownership to the original primary. Protection status of the VM shows as **Protected**.

Protecting Virtual Machines After Disaster

In the event of a disaster, you may lose the source site altogether. After the recovery is performed, you may want to protect the recovered VMs to a newer cluster.

-
- Step 1** Recover the Virtual Machines. Perform standalone recovery (Recovering VMs) or group recovery (Recovering VMs in protection groups). See [Recovering Virtual Machines, on page 30](#) for more details.
 - Step 2** To clear the existing pairing, run the following command in the HX Connect WebCLI:


```
stcli dp peer forget --all
```

 Now the cluster is no longer paired to the original source.
 - Step 3** Unprotect all the local and remote VMs. See [Unprotecting Virtual Machines, on page 27](#) for more details.
 - Step 4** Pair to the new cluster. See the [Creating a Replication Pair, on page 17](#) section for more details.
 - Step 5** Protect the virtual machines.
-

Deleting Protected Virtual Machines

To delete protected virtual machines using the HX Data Platform Command Line Interface (CLI) interface, do the following:

-
- Step 1** Identify the name of the virtual machine that has to be deleted.
 - Step 2** Log in to the controller VM of the virtual machine.
 - Step 3** Run the command `stcli dp VM list --brief` to identify the virtual machine ID for that virtual machine.

This example shows a summary list of the protected virtual machines.

```
stcli dp vm list --brief
```

```
vmInfo:
```

```
-----
name: HC-Edge1-VM30
uuid: 4224f065-2fd9-e4d6-242e-b7a272202b38
-----
name: HC-Edge1-VM1
uuid: 422417f9-c902-59fe-4e9d-8298c2517410
-----
name: HC-Edge1-VM28
uuid: 4224cfb4-5a78-dced-ca5f-3dd6e132492e
-----
name: HC-Edge1-VM13
uuid: 42249800-035b-0c92-0c5f-dcfbc9a5b9e
-----
```

- Step 4** Run the command `stcli dp vm delete --vmid <vmid>` to delete the required virtual machines.

This example shows how to delete a virtual machine.

```
stcli dp vm delete --vmid 4224f065-2fd9-e4d6-242e-b7a272202b38
```

Replication Maintenance Overview

Replication, when configured, runs in the background according to defined schedules. Replication maintenance tasks include:

- **Testing recovery**—Testing if the recovery methods are working. See [Testing Virtual Machine Recovery, on page 29](#) for more details.

- **Pausing replication**—When preparing to upgrade the HX cluster and replication is configured, replication activity must be paused.

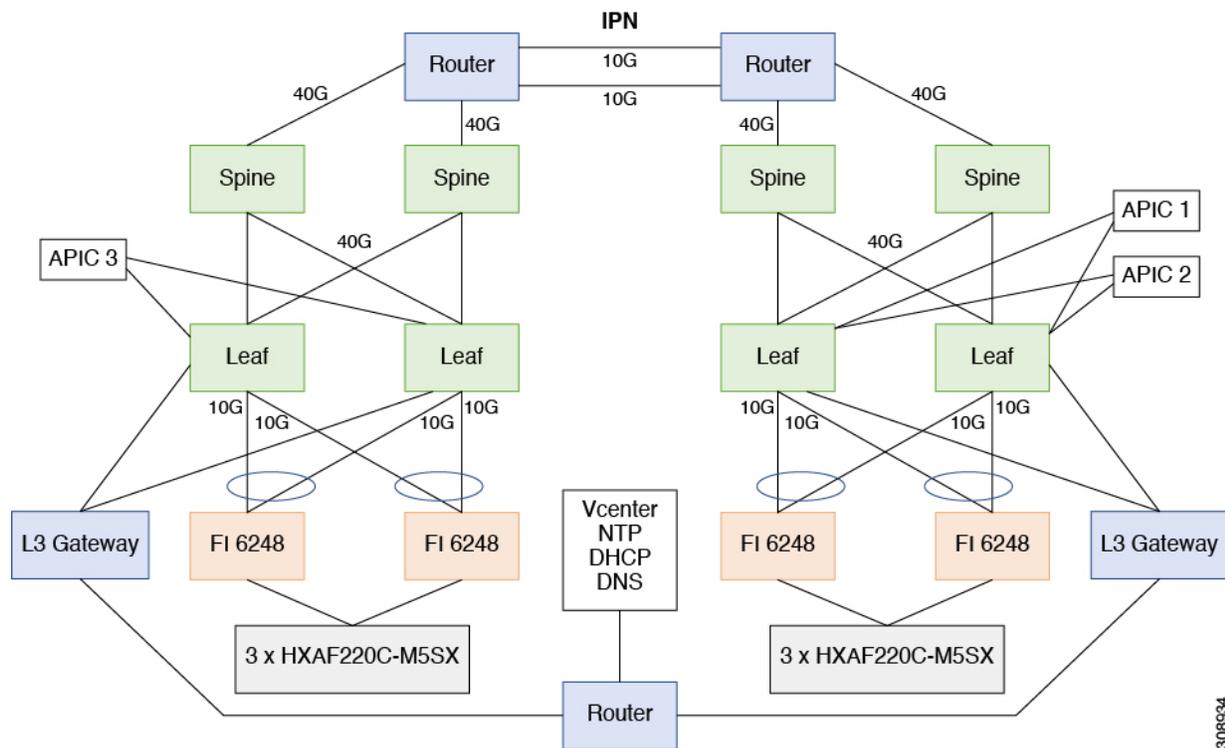
Use the `stcli dp schedule pause` command.

- **Resuming replication**—After HX cluster maintenance activities are complete, resume the replication schedule.

Use the `stcli dp schedule resume` command.

- **Migration**—The option to shift VMs from one source cluster to the replication paired target cluster, making the target cluster the new source cluster for the migrated VMs.

The following image illustrates which configuration is used for Disaster Recovery on HyperFlex if you are deploying in an ACI setup on a broad scale:



308934

Pausing Replication

Before performing a storfs or platform upgrade, if replication is configured replication activity must be paused.

-
- Step 1** Log in to a Storage Controller VM.
 - Step 2** From the command line, run the `stcli dp schedule pause` command.
 - Step 3** Perform the upgrade task.
 - Step 4** Resume the replication schedule.
-

Resuming Replication

After successfully upgrading the HX Storage Cluster which had replication configured, do the following to resume the replication schedule.

Before you begin

Ensure that HX cluster replication is paused and that any maintenance or upgrade tasks are complete.

-
- Step 1** Log in to a Storage Controller VM.
 - Step 2** From the command line, run the `stcli dp schedule resume` command.
-

The previously configured replication schedule for all the protected virtual machines begins.