



Troubleshooting

- [Viewing Resiliency Status in HX Connect, on page 1](#)
- [Troubleshooting the Network Configuration, on page 3](#)
- [Increasing Datastore Capacity May Not Show as Free Space, on page 3](#)
- [Troubleshooting for Site-to-Site Failover, on page 4](#)

Viewing Resiliency Status in HX Connect

The Dashboard page in HX Connect displays the status summary of your HX storage cluster. Click **Information** (i) to access the resiliency status. Resiliency status is the data resiliency health status and the ability of the storage cluster to tolerate failures.

To access the resiliency status in HX Connect:

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. Click **Login**.
2. On the left navigation pane, click **Dashboard**.
3. On the **Dashboard** page, you can view the resiliency status for the HyperFlex Stretched Cluster:

Color coding and icons are used to indicate various status states. Click an icon to display additional information, such as reason messages that explain what is contributing to the current state. The various Resiliency Status states are:

- **Healthy**—The cluster is healthy with respect to data and availability.
- **Warning**—Either the data or the cluster availability is being adversely affected.
- **Unknown**—A transitional state while the cluster is coming online.



Important

The disk failures that are mentioned below are for the converged nodes only (the witness node and the converged nodes do not have disks hosting user data).

Independent and Non-Simultaneous Failure Scenarios

Failure Scenario	Expected Behavior	Resiliency Status in HX Connect
One node failure	VMs will failover to the remaining nodes on the same site as long as those nodes can accommodate the VM resources.	Warning—Cluster shows status of unhealthy until the cluster recovers.
All node failure on any one site	User VMs failover to other site. Note For the recovery operation to succeed after a site failure, the witness node must be online.	Warning—HX Connect shows the site failure details. Cluster shows status of unhealthy until the cluster recovers.
One disk failure	Cluster recovers after failover.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Two disk failure on a single site (one disk on each node) simultaneously	Same as one disk failure.	Warning—Cluster shows status of unhealthy until the cluster recovers. Same as one disk failure.
Witness failure	Cluster remains online.	No visible indication.
vCenter failure (platform impact)	Cluster remains online.	No visible indication.
Network isolation between a given site and a witness	Cluster remains online.	No visible indication.
Network isolation between sites	All user VMs from one site will fail over to the other site. The VMs from the other site will continue to run. Note For the recovery operation to succeed after a site failure, the witness node must be online.	Warning—HX Connect shows the site failure details. Cluster shows status of unhealthy until the cluster recovers.

Failure Scenarios with Witness Failure

Failure Scenario	Expected Behavior	Visible Indications in HX Connect
Disk failure	Cluster remains online.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Single node failure	Cluster remains online. Cluster becomes unhealthy.	Warning—Cluster shows status of unhealthy until the cluster recovers.

Failure Scenario	Expected Behavior	Visible Indications in HX Connect
Link failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness is restored.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Either fabric interconnect pair failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness is restored.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Switch failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness is restored.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Site power failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness is restored.	Warning—Cluster shows status of unhealthy until the cluster recovers.

Troubleshooting the Network Configuration

Do the following to verify if the network configuration prerequisites are met:

- Ensure that the nodes physically residing on different sites are pingable, once the stretched VLANs are created.
- Ensure that the Round-Trip Time (RTT) and bandwidth requirements are met across all the sites. You can use iPerf and Ping or any Cisco supported tool to measure the RTT and bandwidth.

Increasing Datastore Capacity May Not Show as Free Space

Issue:

Increasing the datastore capacity may not show as free space in the HX Connect UI and stcli.

Recommended Solution:

Resize the datastore depending on the actual space that is used by each VM.



Note Use the maximum cluster capacity to resize the datastore. It reflects the actual used spaced of the datastore, which can be used as a hint to resize the datastore based on your requirements.

1. From HX Connect, select **Datastores**.
 2. Select a *datastore*. Click **Edit**.
 3. In the **Edit Datastore** dialog box, to resize the datastore, change the **Size** as needed.
 4. Click **Edit Datastore**.
-

Troubleshooting for Site-to-Site Failover

Do the following during site-to-site failover (if needed):

- During site-to-site failure, depending on the time it takes to failover, VMs may see a transient All Paths Down (APD) event from the ESX layer. This APD is expected when the witness connection is slow.
- In some cases, you may need to manually restart the user VMs if they have timed-out during site-to-site failover.

Do not try to resolve the APD by rebooting the nodes. Depending on the size of the cluster and amount of data, the failover can take 20-30 minutes or more if the witness transaction latencies are high.