




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 4](#)
- [Troubleshooting for Site-to-Site Failover, on page 5](#)
- [Redeploy HyperFlex Stretch Cluster Witness VM, on page 5](#)

Viewing Resiliency Status in HX Connect

The Dashboard page in HX Connect displays the status summary of your HX storage cluster. Click **Information**

() 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 into 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 Stretch 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 VM or Invisible Cloud Witness node and the compute 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 VM or Invisible Cloud 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 VM or Invisible Cloud Witness failure	Cluster remains online.	Witness VM :No visible indication. Invisible Cloud Witness: Online/Offline/Unknown is shown on the System Information Page
vCenter failure (platform impact)	Cluster remains online.	No visible indication.
Network isolation between a given site and a witness VM or Invisible Cloud 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 VM or Invisible Cloud 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 VM or Invisible Cloud 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.
Link failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness VM or Invisible Cloud 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 VM or Invisible Cloud 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 VM or Invisible Cloud 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 VM or Invisible Cloud 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 stretch 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 VM or Invisible Cloud 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 VM or Invisible Cloud Witness transaction latencies are high.

Redeploy HyperFlex Stretch Cluster Witness VM

In some circumstances, it may be required to re-deploy the Stretch cluster witness VM.

Some examples include:

- Witness VM is deleted
- Witness VM is corrupted



Note This process should not be used in an attempt to change IP of the Witness VM.

To redeploy the Stretch Cluster Witness VM:

- Step 1** Download the appropriate version of Stretch cluster witness .ova.
- Step 2** Verify the cluster is healthy.
- Step 3** Deploy a new witness VM with the same IP address as the previous one.
- Step 4** Log into the new witness VM and stop exhibitor.

Example:

```
root@Cisco-HX-Witness-Appliance:~# service exhibitor stop
```

- Step 5** Locate an SCVM in the cluster which is part of the zookeeper ensemble.

Note There should be two nodes per site that are part of the ensemble, and they show as either **leader** or **follower** in the example output.

Example:

```
root@SCVM-1:~# echo srvr | nc localhost 2181
Zookeeper version: 3.4.6--1, built on 07/26/2017 20:05 GMT
Latency min/avg/max: 0/0/3043
Received: 531158372
Sent: 531194128
```

```

Connections: 9
Outstanding: 0
Zxid: 0x350002e75c
Mode: follower <<<<<<<<<<<<
Node count: 3804

```

Step 6 Copy the exhibitor.properties file from the SCVM to the witness VM under the /tmp location.

Example:

```

root@SCVM-1:~# scp /usr/share/exhibitor/exhibitor.properties
root@<Witness-VM-IP>:/tmp/exhibitor.properties

```

Step 7 Edit the file on the witness VM and locate the line starting with **com.netflix.exhibitor.servers-spec**.

Example:

```

root@Cisco-HX-Witness-Appliance:~# vi /tmp/exhibitor.properties#

Auto-generated by Exhibitor
#Mon Oct 29 15:51:29 PDT 2018
com.netflix.exhibitor-rolling-hostnames=
com.netflix.exhibitor-rolling.zookeeper-data-directory=/var/zookeeper
com.netflix.exhibitor-rolling.servers-spec=0\:192.168.4.159,1\:10.9.47.50
,2\:10.9.47.49,3\:10.9.47.45,4\:10.9.47.44
com.netflix.exhibitor.java-environment=ZOO_LOG_DIR=/var/log/zookeeper\
nZOO_LOG4J_PROP="INFO,ROLLINGFILE"nZOOPIDFILE=/tmp/zookeeper/
zookeeper_server.pid
com.netflix.exhibitor.zookeeper-data-directory=/var/zookeeper
com.netflix.exhibitor-rolling-hostnames-index=0
com.netflix.exhibitor-rolling.java-environment=ZOO_LOG_DIR=/var/log/
zookeepernZOO_LOG4J_PROP="INFO,ROLLINGFILE"nZOOPIDFILE=/tmp/zookeeper/
zookeeper_server.pid
com.netflix.exhibitor-rolling.observer-threshold=0
com.netflix.exhibitor.servers-spec=0\:192.168.4.159,1\:10.9.47.50,2\
:10.9.47.49,3\:10.9.47.45,4\:10.9.47.44
com.netflix.exhibitor.cleanup-period-ms=0
com.netflix.exhibitor.auto-manage-instances-fixed-ensemble-size=0
com.netflix.exhibitor.zookeeper-install-directory=/usr/share/zookeeper
com.netflix.exhibitor.check-ms=30000

```

Note In the above example, the 10.9.47.x IP's are the storage data subnet while the 192.168.4.x IP's are the management subnet.

Step 8 With all of the data IP addresses of SCVM's in the zookeeper cluster, and the management IP address of the witness VM displayed, replace all of the data IP addresses with the management IP addresses for each corresponding controller VM.

Example:

```

root@Cisco-HX-Witness-Appliance:~# vi /tmp/exhibitor.properties#

Auto-generated by Exhibitor
#Mon Oct 29 15:51:29 PDT 2018
com.netflix.exhibitor-rolling-hostnames=
com.netflix.exhibitor-rolling.zookeeper-data-directory=/var/zookeeper
com.netflix.exhibitor-rolling.servers-spec=0\:192.168.4.159,1\:10.9.47.50,2\:10.9.47.49,3\:10.9.47.45,4\:10.9.47.44
com.netflix.exhibitor.java-environment=ZOO_LOG_DIR=/var/log/zookeepernZOO_LOG4J_PROP="INFO,ROLLINGFILE"nZOOPIDFILE=/tmp/zookeeper/zookeeper_server.pid
com.netflix.exhibitor.zookeeper-data-directory=/var/zookeeper
com.netflix.exhibitor-rolling-hostnames-index=0
com.netflix.exhibitor-rolling.java-environment=ZOO_LOG_DIR=/var/log/zookeepernZOO_LOG4J_PROP="INFO,ROLLINGFILE"nZOOPIDFILE=/tmp/zookeeper/zookeeper_server.pid
com.netflix.exhibitor-rolling.observer-threshold=0
com.netflix.exhibitor.servers-spec=0\:192.168.4.159,1\:192.168.4.50,2\:192.168.4.49,3\:192.168.4.45,4\:192.168.4.44
com.netflix.exhibitor.cleanup-period-ms=0
com.netflix.exhibitor.auto-manage-instances-fixed-ensemble-size=0

```

```
com.netflix.exhibitor.zookeeper-install-directory=/usr/share/zookeeper  
com.netflix.exhibitor.check-ms=30000
```

Step 9 Save changes to the file and copy the file from /tmp to /usr/share/exhibitor/.

Example:

```
root@Cisco-HX-Witness-Appliance:~# cp /tmp/exhibitor.properties  
/usr/share/exhibitor/exhibitor.properties
```

Step 10 Wait for the witness VM to sync with the cluster (this is typically quick but can take up to an hour). You can verify the Witness VM is online from HX Connect under **System Information**.
