



Preinstallation Checklist

- [Invisible Cloud Witness Preinstallation Checklist for All New Installs, on page 1](#)
- [Common Site Interlink and Witness Requirements, on page 2](#)

Invisible Cloud Witness Preinstallation Checklist for All New Installs

This checklist applies to all new Invisible Cloud Witness Stretch Cluster installs using HXDP 5.5(1a) and later.



Remember All new Stretch Cluster installs auto-configure an Invisible Cloud Witness for site arbitration. Fresh installs using the Witness VM (HXDP 5.0(x) and earlier) is not supported.

- Network connectivity for the install powered by Installer VM.



Note Intersight deployment of stretch-cluster is NOT supported

- Intersight access
- Auxiliary ZooKeeper (AUX ZK IP): This must be in the same data network.
- Preferred site: The site designated to serve requests in the event that the connectivity between sites is down.



Note If you have an existing Stretch Cluster (HXDP Release 5.0(x) and earlier), you will continue to use upgrade the VM based witness as described in the [Cisco HyperFlex Systems Stretch Cluster Guide, Release 5.0](#) and the [Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Release 5.0](#).

Common Site Interlink and Witness Requirements

The following describe general requirements, and the HXDP release they apply to.

Network Requirements

- 10 Gbps dedicated, 5-ms Round-Trip Time (RTT) latency between the two active sites is required.
- 100 Mbps, 200 ms-RTT worst case latency for 16 kilobyte packet sizes between the active sites and witness site is required.
- Existing fabric interconnects are supported, provided the fabric interconnects support M5 servers.
- User VMs should be capable of vMotioning to any site, without impacting external network connectivity to these VMs.
- Similar to regular HX Clusters, two separate IP subnets are required—Both over Stretch L2. One subnet for data traffic and one for management traffic, with the management subnet reachable from vCenter and the witness node.
- Static IP address assigned to Witness VM, can be changed only with cluster redeployment. If DHCP server is used to define the network configuration, the IP address needs to be dedicated to Witness VM.
- FI facing ports need to have Port-fast, spanning-tree port type edge trunk, or similar spanning tree configuration that immediately put ports into forwarding mode.
- QoS


If the HyperFlex cluster is a stretched cluster, you should enable QoS end-to-end between the sites. The QoS policies in the Cisco UCS and upstream network should also be aligned so that HyperFlex traffic can receive consistent QoS end-to-end

If the HyperFlex cluster is a stretched cluster that connects to an ACI Multi-Pod fabric, you should enable QoS policies in the ACI fabric and extend it across the IPN as well.

The QoS policies in the Cisco UCS and ACI fabrics should also be aligned so that HyperFlex traffic can receive consistent QoS end-to-end.

There may be a potential conflict between the Quality of Service (QoS) policy for HyperFlex and Nexus 1000v. Make sure that the QoS classes for N1Kv are set as per the HyperFlex policy. See [Creating a QoS Policy](#), in the [Network and Storage Management Guide](#).

Network Topology



Witness VM Requirements

- Configure and enable NTP on all servers.
- An independent third witness site is required.
- A separate witness is required for each cluster.
- Both the main sites must have connectivity to the third witness site with a minimum bandwidth of 100 MBPS, 200 ms-RTT worst case latency for 16 kilobyte packet sizes.
- Site must have the capability to deploy and run Open Virtualization Format (OVF) image.

Network Latency Requirements for the Witness VM

- The HyperFlex Stretched Cluster solution requires that the Witness VM be located in a third site to ensure that a storage site failure does not affect the Witness VM.
- The solution can support a Witness bandwidth as low as 100 Mbps, 200 ms-RTT worst case latency for 16 kilobyte packet sizes.
- Latency to the witness impacts site failure times, and it is recommended that for larger clusters with significant load and data, to have RTT times in the order of 10ms or lower.

Witness VM Fabric Interconnect and Node Requirements

- Symmetric configuration is required across both sites.
- There must be a minimum of two converged nodes on each site.

- A maximum of 16 converged nodes on each site is supported. Ensure that both sites have the same number of converged nodes.
- There must be a redundant fabric interconnect configuration on each site.
- Converged nodes have to be M5 or M6 nodes.
- Ensure that the Fabric Interconnect pair is of the same model in the same domain.
- Compute-only nodes are supported.

VMware Requirements

- VMware Enterprise Plus edition with HyperFlex Stretched Cluster is highly recommended to ensure proper failover behavior and guarantee high performance during normal operations. While it is possible to run Stretched Cluster without VMware Enterprise Plus edition, the advanced DRS features such as site affinity will not be available, negating some of the intended operational functions of Stretched Cluster.
- Use a single vCenter for both sites.
- The vCenter can be a VM running at the same site as the witness.
- Nested vCenter is not supported in Stretched Cluster.
- The vCenter must be configured independently for High Availability, as required.

UCS Manager Requirements

- Two separate, manually synchronized fabric interconnect domains are required.

VLAN Requirements

- IP addresses for nodes on both sites are required.
- Stretched VLANs across both sites are required.

Stretch Witness VM

Required only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.

Table 1: Port Requirements

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
2181 2888 3888	(Zookeeper lifecycle)/TCP	Witness	Each CVM Node	Bidirectional, management addresses
8180	Exhibitor (Zookeeper lifecycle)/TCP	Witness	Each CVM Node	Bidirectional, management addresses
80	HTTP/TCP	Witness	Each CVM Node	Potential future requirement

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
443	HTTPS/TCP	Witness	Each CVM Node	Potential future requirement

Deploying the Witness VM Node



Important

- A HyperFlex Witness VM node is mandatory in a stretch cluster environment to achieve quorum in case of total failure in any of the sites or when the network link between the sites encounters a failure.
- In case of accidental deletion or loss of Witness VM, to replace the Witness VM please contact Cisco TAC.
- The Witness VM requires a Static IP address that cannot be changed without cluster redeployment. If DHCP server is used to define the network configuration, the IP address needs to be dedicated to Witness VM.



Note

If the IP address is requested from a DHCP server please verify that the clustering service is running after deploying the Witness VM, perform the following steps:

- Check for the clustering service by running the following command after every reboot of the Witness VM:


```
service exhibitor status
```
- If the output does not show the service as running, then restart exhibitor using the following command after ensuring that the witness VM does have an IP address:


```
service exhibitor restart
```

The following procedure details the steps to follow, for deploying a witness VM virtual node on a physical ESXi host.

Before you begin

This section is for use only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.



Attention

- HyperFlex witness VM node version 1.1.3 is supported in Cisco HXDP Release 4.5(2a) and later. To review the recommended version for your specific release, see the *HX Data Platform Software Versions for HyperFlex Witness Node* sections of the [Cisco HyperFlex Software Requirements and Recommendations](#) guide.
- The Witness VM must be deployed on an ESXi server that has sufficient hosting capability. The Witness VM requires 4 vCPUs, 8GB of memory, and 40 GB of disk space.
- Ensure that the virtual network on this ESXi host is reachable from both the stretch cluster sites.

- Download the HyperFlex witness VM node on to your desktop or host that runs vSphere Web Client from [Download Software](#).

Example:
HyperFlex-Witness-1.0.2.ova

- High Availability is optional for the witness VM node.

Step 1 Log into vSphere Web Client. Choose the ESXi server where the witness VM should be deployed. Right-click the ESXi host and select **Deploy OVF Template**.

Step 2 Browse and select the *HyperFlex-Witness.ova* file. Click **Next**.

Step 3 Specify a unique name for the witness VM node in the **Virtual Machine Name** field. Select a location for the virtual machine from the drop-down list. Click **Next**.

Step 4 From the **Select a compute resource** drop-down list, choose the ESXi host where you want to deploy the witness VM node. Click **Next**.

Step 5 In the **Review details** pane, verify the template details. Click **Next**.

Step 6 In the **Select Storage** pane, do the following:

Field	Description
Select virtual disk format drop-down list	<ul style="list-style-type: none"> • Thick Provision Lazy Zeroed • Thick Provision Eager Zeroed • Thin Provision
VM Storage Policy drop-down list	Datastore Default

Select the datastore where the virtual machine will reside. Ensure that this datastore has at least 40 GB of available free space. Click **Next**.

Step 7 In the **Select Networks** pane, select a *Destination Network* port group, where the witness VM has to connect. Click **Next**.

Step 8 On the **Customize Template** page, complete the fields that are relevant for your configuration. If no values are entered, the VM uses DHCP server provided network configuration parameters.

Field	Description
Static IP Address field	<p>The IP address for Witness VM.</p> <p>The Static IP address can be changed only with cluster redeployment. If the DHCP server is used to define the network configuration, the IP address needs to be dedicated to Witness VM.</p> <p>Leave blank if DHCP is desired.</p>
Netmask field	<p>The netmask or prefix for this interface.</p> <p>Leave blank if DHCP is desired.</p>
Default Gateway field	<p>The default gateway address for this VM.</p> <p>Leave blank if DHCP is desired.</p>

Field	Description
DNS field	The domain name servers for this VM (comma separated). Leave blank if DHCP is desired.
NTP field	NTP servers for this VM (comma separated) to sync time. Leave blank if DHCP is desired.

Click **Next**.

Step 9 On the **Ready to complete** page, verify all the details entered. Click **Finish**.

Step 10 Repeat this process for each cluster.

What to do next

After successfully deploying the witness node, you can proceed to installing your Stretch cluster. When prompted enter the IP address of this witness node on the **IP Address** page when creating a HyperFlex Stretch cluster. The witness node is automatically used during configuration of the stretch cluster.

Changing the Witness VM Password

After successful deployment of witness VM, you must change the default password.

Before you begin

Required only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.

Download and deploy the witness VM.

Step 1 Log into the witness VM using SSH.

```
$ ssh root@<IP address of witness VM>
```

Step 2 Enter the default password.

Step 3 Enter the `passwd` command and change the password.

```
$ passwd
Changing password for user admin.
```

Step 4 Log out from the witness VM.

What to do next

Log into the witness VM with the new password.

Witness VM Network IP Addressing

Required only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.

IP addresses for HyperFlex Stretch Cluster need to be allocated from the appropriate subnets and VLANs to be used.

**Important**

- Required only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.
- Ensure that the Data and Management Networks are on different subnets for a successful installation.
- Ensure that the IP addressing is in the same subnet for the same components (Management Network IP addresses, Data Network IP addresses).

Table 2: Example: Stretch Cluster IP Addressing

Storage Cluster Management IP address	10.10.10.128	Storage Cluster Data IP address	192.168.10.160
Subnet mask IP address	255.255.255.0	Subnet mask IP address	255.255.255.0
Default gateway IP address	10.10.10.1	Default gateway IP address	192.168.10.1

Table 3: Example: Stretch Cluster IP Addressing for Site A

Management Network IP Addresses (must be routable)			Data Network IP Addresses (does not have to be routable)	
ESXi Hostname*	Hypervisor Management Network	Storage Controller Management Network	Hypervisor Data Network (Not Required for Cisco Intersight)	Storage Controller Data Network (Not Required for Cisco Intersight)
Server 1:	10.10.10.2	10.10.10.32	192.168.10.2	192.168.10.32
Server 2:	10.10.10.3	10.10.10.33	192.168.10.3	192.168.10.33
Server 3:	10.10.10.4	10.10.10.34	192.168.10.4	192.168.10.3
Server 4:	10.10.10.5	10.10.10.35	192.168.10.5	192.168.10.35
Server 5:	10.10.10.6	10.10.10.36	192.168.10.6	192.168.10.36

Table 4: Example: Stretch Cluster IP Addressing for Site B

Management Network IP Addresses (must be routable)			Data Network IP Addresses (does not have to be routable)	
ESXi Hostname*	Hypervisor Management Network	Storage Controller Management Network	Hypervisor Data Network (Not Required for Cisco Intersight)	Storage Controller Data Network (Not Required for Cisco Intersight)
Server 1:	10.10.10.64	10.10.10.96	192.168.10.64	192.168.10.96

Management Network IP Addresses (must be routable)			Data Network IP Addresses (does not have to be routable)	
Server 2:	10.10.10.65	10.10.10.97	192.168.10.65	192.168.10.97
Server 3:	10.10.10.66	10.10.10.98	192.168.10.66	192.168.10.98
Server 4:	10.10.10.67	10.10.10.99	192.168.10.67	192.168.10.99
Server 5:	10.10.10.68	10.10.10.100	192.168.10.68	192.168.10.100

