



# Preinstallation Checklist

- [Common Site Interlink and Witness Requirements, on page 1](#)
- [Invisible Cloud Witness Preinstallation Checklist for All New Installs, on page 6](#)
- [Stretch Witness VM, on page 6](#)
- [Using the Intersight Arbitrator Swichover , on page 7](#)
- [Changing the Witness VM Password, on page 10](#)
- [Witness VM Network IP Addressing, on page 11](#)

## Common Site Interlink and Witness Requirements


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

### Network Requirements

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
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 or M6 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.	✓	✓

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (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.	✓	✓
<p>QoS</p> <p>If the HyperFlex cluster is a stretch 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.</p> <p>If the HyperFlex cluster is a stretch 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.</p> <p>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.</p> <p>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 <a href="#">Creating a QoS Policy</a>, in the <a href="#">Network and Storage Management Guide</a>.</p>	✓	✓

### Network Topology



**Witness VM Requirements**

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
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**

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
The HyperFlex Stretch 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**

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
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**

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
VMware Enterprise Plus edition with HyperFlex Stretch Cluster is highly recommended to ensure proper failover behavior and guarantee high performance during normal operations. While it is possible to run Stretch 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 Stretch 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 Stretch Cluster.	✓	✓
The vCenter must be configured independently for High Availability, as required.	✓	✓

**UCS Manager Requirements**

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
Two separate, manually synchronized fabric interconnect domains are required.	✓	✓

**VLAN Requirements**

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
IP addresses for nodes on both sites are required.	✓	✓
Stretch VLANs across both sites are required.	✓	✓

# 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.

- Using the Intersight Arbitrator instead of witness setup for switchover is supported.
- Network connectivity for the install powered by Installer VM.
- 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](#).

## 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
443	HTTPS/TCP	Witness	Each CVM Node	Potential future requirement

## Using the Intersight Arbitrator Swichover

Beginning with HXDP Release 6.0(x), you can use the Intersight Arbitrator instead of witness setup for swichover



- 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.

Complete the following tasks to deploy Stretched Cluster Witness using Intersight's Arbitrator for swichover using `hxcli`.

### SUMMARY STEPS

1. Run the `hxcli` command for swichover from Controller vm.
2. To check the current witness related details of the stretch cluster use the command `hxcli cluster witness`.
3. To locate the initiate swichover parameters definitions, use the command `hxcli cluster swichover -h`.
4. To initiate swichover, use the command `hxcli cluster swichover`.

- Review your Stretch Cluster Infrastructure using the command `hxcli cluster witness`.

## DETAILED STEPS

- Step 1** Run the `hxcli` command for swichover from Controller vm.

```

hxshell:~$ hxcli cluster
Displays the list of commands available in the cluster namespace.
Usage :
hxcli cluster [flags]
hxcli cluster [command]
Aliases :
cluster, cl
Available Commands:
about Cluster version, model and other metadata details
detail Short summary of cluster details
health Health and resiliency information of the cluster
info Detailed information about cluster
shutdown Shuts down the Hyperflex cluster
start Start the Hyperflex cluster
stats Cluster capacity and space savings
Cluster Operational Status and Resiliency Health
Stretch Cluster witness infrasturcture swichover process
Stretch Cluster witness infrastructure status
Flags:
-h, -help help for cluster
status
swichover
witness

```

- Step 2** To check the current witness related details of the stretch cluster use the command `hxcli cluster witness`. The following example shows the output of command before swichover is done on a Stretch cluster. Note that the Stretch Cluster has `ZK_WITNESS` and the witness ip is displayed to user.

```

hxshell:~$ hxcli cluster witness
Stretch Cluster witness type : ZK_WITNESS Stretch Cluster witness ip : 10.198.148.26
Stretch Cluster witness state : ONLINE hxshell:~$ |

```

- Step 3** To locate the initiate swichover parameters definitions, use the command `hxcli cluster swichover -h`.

```

hxshell:~$ hxcli cluster swichover -h
Stretch Cluster witness infrasturcture swichover process

Usage:
  hxcli cluster swichover [flags]

Aliases:
  swichover, so

Flags:
  -l, --arb-ip string      Stretch Cluster witness local arbitrator ip
  -n, --arb-name string    Stretch Cluster witness local arbitrator user name
  -p, --arb-pwd string     Stretch Cluster witness local arbitrator password
  -t, --arb-type string    Stretch Cluster witness arbitrator type [intersight | local]
  -a, --aux-ip string     Stretch Cluster witness aux ip
  -h, --help              help for swichover
  -s, --site string       Stretch Cluster witness preferred site

```

### Common Error Messages: Input Arguments Checking



- a) **Aux IP Expectations:** Aux ip provided if not valid, unused, and present in same subnet of the data network, is reported as invalid input with the following error message:

```
Internal Server Error (500): ("cause":"Aux ip provided should be valid, unused, and present in same subnet as the data network ", "errorCode":400, "message":"!)
```

- b) **Site Input Expectations:** Site provided should be one of your stretch cluster sites. If not, it is reported as invalid parameter with the following error message:

```
Internal Server Error (500): ("cause":"Field name: site has illegal value. Should be one of: SiteA.SiteB", "errorCode":400, "message":"failure: Witness vm to Arbitrator swichover process failed.")
```

- c) **Arbitrator Type Expectations:** Arbitrator type should be local. If not it is reported as an invalid parameter with the following error message:

```
Internal Server Error (500): ("cause":"Field name: arb type has illegal value. Should be one of: intersight local", "errorCode":400, "message":"failure: Wittness vm to Arbitrator swichover process failed.")
```

#### Step 4

To initiate swichover , use the command `hxcli cluster swichover`.

The following example shows a successful swichover.

```
hxshell:~$ hxcli cluster swichover --aux-ip 172.16.60.102 --site SiteA --arb-type intersight

Stretch cluster swichover from Witness vm to Arbitrator has been initiated..
Swichover process might take 5 - 6 mins or more based on certain conditions: /

Awaiting ensemble changes to update on all nodes
Swichover process might take 5 - 6 mins or more based on certain conditions: / Task Id : 750975cb-abd6-433a-aa5d-1c419e4ee4a2
Description : Notify swichover ensemble changes to all
              : nodes
Status      : Succeeded
Name       : Swichover ensemble update workflow
Category   : STRETCH_CLUSTER_SWICHOVER
Submitted At : 2024-04-05 01:06:28.804 -0700 PDT
Swichover process might take 5 - 6 mins or more based on certain conditions: -
Initiating conversion of standalone nodes to observers

Stretch Cluster witness swichover to arbitrator success

Stretch Cluster preferred site has been updated successfully
hxshell:~$
```

**Note** Post swichover, standalone nodes are converted to observers. To verify conversion, use the command `zkstatus`.

#### Step 5

Review your Stretch Cluster Infrastructure using the command `hxcli cluster witness`.

Example output:

```
hxshell:~$ hxcli cluster witness
Stretch Cluster witness type : INTERSIGHT
Stretch Cluster witness ip : INTERSIGHT
Stretch Cluster witness state : ONLINE
```

System Overview Nodes Disks Last refreshed at: 04/05/2024 1:43:54 PM

Cluster not registered with Cisco Licensing. [Register Now](#)

**simulator\_witness**  
ONLINE

License Type Evaluation  
License Status License expires in 89 days. Cluster not registered with Cisco Licensing.

vCenter https://10.198.106.14  
Uptime 0 days, 20 hours, 49 minutes, 48 seconds

Hypervisor 7.0.3-20842708  
HXDP Version 6.0.1a-44239  
Witness **Online (INTERSIGHT)**

Total Capacity 1.84 TiB  
Available Capacity 1.8 TiB  
Data Replication Factor 2+2

DNS Server(s) 10.107.128.8  
NTP Server(s) 10.107.129.14  
Controller Access over SSH Enabled

SiteA **PREFERRED SITE** | Hyperconverged Nodes

Disk View Options [Disk View Legend](#)

---

HyperFlex Connect simulator\_witness

Dashboard

MONITOR

Alarms

Events

**Activity**

ANALYZE

Performance

PROTECT

Notify switchover ensemble changes to all nodes  
Status: Success  
04/05/2024 1:36:28 PM

Switch a stretch cluster from using Witness VM to a local/intersight arbitrator  
Status: Success  
04/05/2024 1:36:23 PM

Validate inputs and cluster condition  Validate inputs and cluster condition

Remove Witness VM and Reconfigure Storage Cluster  Remove Witness VM and Reconfigure Storage Cluster

Configuring new services to provide Arbitration Capability  Configuring new services to provide Arbitration Capability

### 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**

<b>Storage Cluster Management IP address</b>	10.10.10.128	<b>Storage Cluster Data IP address</b>	192.168.10.160
<b>Subnet mask IP address</b>	255.255.255.0	<b>Subnet mask IP address</b>	255.255.255.0
<b>Default gateway IP address</b>	10.10.10.1	<b>Default gateway IP address</b>	192.168.10.1

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

<b>Management Network IP Addresses (must be routable)</b>			<b>Data Network IP Addresses (does not have to be routable)</b>	
<b>ESXi Hostname*</b>	<b>Hypervisor Management Network</b>	<b>Storage Controller Management Network</b>	<b>Hypervisor Data Network (Not Required for Cisco Intersight)</b>	<b>Storage Controller Data Network (Not Required for Cisco Intersight)</b>
Server 1:	10.10.10.2	10.10.10.32	192.168.10.2	192.168.10.32

<b>Management Network IP Addresses (must be routable)</b>			<b>Data Network IP Addresses (does not have to be routable)</b>	
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*

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