



Guidelines and Limitations

- [Guidelines, on page 1](#)
- [Limitations, on page 2](#)

Guidelines

Consider the following guidelines when creating a HyperFlex Stretch Cluster:

- HXDP Enterprise Edition licensing is required to use the HyperFlex Stretch Cluster feature.
- vSphere Enterprise Plus licensing is required for full featured HyperFlex Stretch Cluster functionalities like VM load balancing and VM migration.
- Ensure that DRS is enabled.
- Supported Stretch Cluster scale:
 - There must be a minimum of 2 nodes on each site. The minimum overall cluster size across both sites is 4.
 - Small Form Factor (SFF). The maximum cluster size across both sites is 64. The maximum converged node count is 16 per site. The compute to converged node ratio can be 2:1, while the maximum limit of nodes per site must not exceed 32. For example, you can have 11 converged and 21 compute nodes per site.
 - Large Form Factor (LFF). The maximum cluster size is 48. The maximum converged node count is 8 per site. The compute to converged node ratio can be 2:1, while the maximum limit of nodes per site must not exceed 24. For example, you can have 8 converged and 16 compute nodes per site.
- There must be a redundant fabric interconnect configuration on each site.
- There must be symmetric cluster configuration across both sites. The number of nodes and the model of HX nodes should be the same on both sites.
- VMs are placed correctly within site affinity only if VMware HA and DRS are enabled before VM creation. Otherwise, the correct placement of VM for proper affinity is not guaranteed.
- HyperFlex Native Replication is supported between Stretch Clusters, and between Stretch Clusters and standard clusters.
- It is a best practice to have two datastores one per site, with the respective affinity.

- **Guest Traffic:** Due to the nature of the Cisco VIC carving up multiple vNICs from the same physical port, it is not possible for guest VM traffic configured on vswitch-hx-vm-network to communicate L2 to interfaces or services running on the same host. It is recommended to either a) use a separate VLAN and perform L3 routing or b) ensure any guest VMs that need access to management interfaces be placed on the vswitch-hx-inband-mgmt vSwitch. In general, guest VMs should not be put on any of the HyperFlex configured vSwitches except for the vm-network vSwitch. An example use case would be if you need to run vCenter on one of the nodes and it requires connectivity to manage the ESXi host it is running on. In this case, use one of the recommendations above to ensure uninterrupted connectivity.

Limitations

Consider the following limitations when creating a HyperFlex Stretch Cluster:

- Self Encrypting Drives (SEDs) are not supported. However, VM based third-party software encryption is supported.
- Overlay networking and L3 protocols are not supported. L2 adjacency is required for the data and management networks.
- Stretch Cluster is not supported on Hyper-V platform.
- Online rolling upgrades are supported only for the HX Data Platform. Cisco UCS Manager upgrades and VMware ESXi upgrades must be performed manually on one node at a time or performed offline.
- Upgrade from standalone cluster to Stretch Cluster configuration is not supported.
- Stretch Cluster is supported only on M5 or M6 nodes. M4/M5, M4/M6, M5/M6 is supported when reached by expansion.
- Shared Witness VM is not supported for Stretch Cluster deployments.