



Cisco HyperFlex Systems Offering Overview

- [About this Guide, on page 1](#)
- [Cisco HyperFlex Systems Offering Overview, on page 1](#)

About this Guide

This publication provides information about ordering Cisco HyperFlex Systems from estimate creation, ordering, contract creation, order holds and expedites, activation, renewals, and co-termination.

Cisco HyperFlex Systems Offering Overview

Cisco HyperFlex HX-Series

Cisco HyperFlex HX-Series provides unified compute, storage, and networking with cloud management. The Cisco HyperFlex HX Data Platform (HXDP) is a high-performance, extensible distributed file system that supports multiple hypervisors with a wide range of enterprise-grade data management and optimization services.

Cisco HyperFlex Clusters

A Cisco HyperFlex cluster is a flexible and highly configurable system built using trusted UCS components. A HyperFlex cluster requires a minimum of three homogeneous nodes (with disk storage) that can scale up to 32 total nodes (refer to the Release Notes documentation for the latest release specific scale support). Data is replicated across at least two of these nodes, and a third node is required for continuous operation in the event of a single-node failure. A pair of Cisco UCS 6400 (10/25Gbps) or 6300 (40Gbps) Series Fabric Interconnects integrate HyperFlex Data Platform (HXDP) licensed Unified Computing System (UCS) nodes into a single hyper-converged system.

HyperFlex clusters are created with homogeneous nodes. All nodes must be the same type and form factor with similar CPU, memory, and drive configurations. The HyperFlex installer does not support mixing node types within the same cluster (e.g. attempting to install 240 nodes and 220 nodes into the same cluster will result in an error message and be blocked by the installer). For more information, see the related installation and expansion documentation.

Cisco HyperFlex provides compute-only nodes, which enable customers to increase CPU, memory, and graphics acceleration capacity without the cost of additional licenses and the ability to scale compute independently of storage. A HyperFlex cluster can contain any combination of three or more HXDP converged

nodes, plus zero or more compute-only nodes, supporting up to a maximum of 32 nodes. The standard HXDP Datacenter Advantage licenses support a 1:1 converged to compute only ratio, or up to the same number of compute-only nodes. Clusters with more compute-only nodes than HXDP-licensed converged nodes require HXDP Datacenter Premier licenses. Please review the HXDP license section of this guide for details on compute-only ratio details and examples.



Note HyperFlex compute-only nodes can also be used as UCS nodes. Cisco will honor the support and warranty terms at the same price. The support and price of the HyperFlex compute-only node hardware is identical to UCS node hardware. HyperFlex converged nodes can be used as UCS nodes. Cisco will honor the support and warranty terms per the original order. Support cost will continue at the HyperFlex rate, even if no longer used as a HyperFlex node. For Hyperflex M4 & M5 : compute nodes can't be used as converged nodes, converged nodes cannot be used as compute-only nodes.

A HyperFlex cluster can reside in a single location or it can be stretched across short geographic distances. An active-active “stretched” cluster synchronously replicates data between the two sites and has a very short recovery time objective (RTO) and zero data loss. Across longer geographic distances, native replication can synchronize data from a primary to a secondary site to support more traditional disaster recovery strategies. HyperFlex also supports stretched clusters where one cluster can span multiple locations. Each location requires a pair of fabric interconnects. For more information about stretched clusters, benefits and configuration considerations, see the [Cisco HyperFlex Systems Stretched Cluster Guide](#).

Cisco HyperFlex Systems M6 Offerings

Cisco HyperFlex M6 Node Configuration Options

You can order HyperFlex M6 nodes using the HX-M6-MLB bundle. This bundle supports the following configurations:

Table 1: HX-M6-MLB Bundle Configuration Options

Type	PID	Description
HyperFlex Hybrid nodes	HX220C-M6S	Cisco HyperFlex Hybrid 220 M6 node
	HX240C-M6SX	Cisco HyperFlex Hybrid 240 M6 node
	HX240C-M6L	Cisco HyperFlex Hybrid 240 M6 node Large Form Factor (LFF)
HyperFlex All Flash nodes	HXAF220C-M6S	Cisco HyperFlex All Flash 220 M6 node
	HXAF220C-M6SN	Cisco HyperFlex All NVMe 220 M6 node
	HXAF240C-M6SX	Cisco HyperFlex All Flash 240 M6 node
	HXAF240C-M6SN	Cisco HyperFlex All NVMe 240 M6 node
HyperFlex Hybrid Edge nodes	HX-E-220M6S	Cisco HyperFlex Hybrid Edge 220 M6 node
	HX-E-240-M6SX	Cisco HyperFlex Hybrid Edge 240 M6 node

Type	PID	Description
HyperFlex All-Flash Edge nodes	HXAF-E-220M6S	Cisco HyperFlex All Flash Edge 220 M6 node
	HXAF-E-240-M6SX	Cisco HyperFlex All Flash Edge 240 M6 node

For more information about HyperFlex nodes, see the [Cisco HyperFlex Technical Specifications](#).

When ordering HX compute-only nodes, use the bundle: **HX-UCSCM6-MLB** (Cisco HyperFlex HX Compute M6 Blade and Rack server MLB).

For more information on software (HXDP) and Intersight licensing, see [M6 Product Information and SKU IDs](#).

Local Storage

The following table lists categories of local storage and how they are utilized by HyperFlex nodes.

Type of Drive	Description
Capacity Drive (Data)	<p>All NVMe nodes use NVMe drives and are the highest performance and lowest latency configuration.</p> <p>All Flash (HXAF) nodes use solid state drives (SSD), which do not have any moving parts, therefore enabling higher performance and reliability.</p> <p>Hybrid (HX) nodes use spinning hard disk drives (HDD) for data storage that provide the lowest cost per Gb of storage.</p> <p>HyperFlex also supports self-encrypting drives (SED) on HXAF and HX nodes. When ordering SED, ensure that the entire HyperFlex cluster utilizes SED encryption.</p> <p>Note HyperFlex does not support mixing SED and non-SED drives in the same cluster.</p>
Cache Drive (Write Log)	<p>Each HX node requires one high-performance drive for data caching and rapid acknowledgment of write requests.</p> <p>All NVMe nodes use Optane 3D Crosspoint drives for caching.</p> <p>All Flash (HXAF) nodes may use SSD or NVMe drives for caching.</p> <p>Hybrid (HX) nodes use SSD drives for caching.</p>
System Drive (Housekeeping)	HX and HXAF nodes use one System Log drive per system.

Drive Support

The following table describes drive support per product.

		System type										
	PID	H200N	H200V	H200S	H200SX	H200MS	H200MX	H200ML	H200S	H200SX	H200S	H200SX
Boot Drive	H200CB	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	H200WRD	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
System Drive (Cache)	H200MV			✓	✓	✓	✓	✓	✓	✓	✓	✓
	H200MH	✓	✓									
Capacity Drive (Data)	H200DN					✓	✓				✓	✓
	H200BN					✓	✓				✓	✓
	H200BN					✓	✓				✓	✓
	H200D9					✓	✓					
	H200D9					✓	✓					
	H200KN						✓					
	H200KN						✓					
	H200KN						✓					
	H200KM						✓					
	H200KM						✓					
	H200KM						✓					
	H200MV			✓	✓				✓	✓		
	H200MV			✓	✓				✓	✓		
	H200MV			✓	✓				✓	✓		
	H200MV			✓	✓				✓	✓		
	H200N9			✓	✓							
	H200N9			✓	✓							
	H200N9			✓	✓							
	H200N9			✓	✓							
	H200N9			✓	✓							

		System type										
	PID	H200N	H200N	H200S	H200X	H200V	H200X	H200L	H200S	H200X	H200S	H200X
	H200N	✓	✓									
	H200N	✓	✓									
	H200N		✓									
Cache Drive (Write Log)	H200N					✓						
	H200N					✓					✓	
	H200N					✓						
	H200N						✓					✓
	H200N						✓					
	H200N							✓				
	H200N			✓	✓				✓	✓		
	H200N			✓								
	H200N			✓								
	H200N	✓	✓		✓							
	H200N				✓							
Drive Controller	H200N			✓	✓	✓	✓		✓	✓	✓	✓
	H200N							✓				



Note Power Cord Compatibility for India

New replacement parts for cords in compliance with the Dec 1st, 2020 Bureau of Indian Standards (BIS) regulatory standard for lower-power consumption products have been added.

For a complete list of part numbers, see the [Cisco HyperFlex Specification Sheets](#).

Cisco HyperFlex Cluster Sizing Guidelines

The Cisco HyperFlex Sizer tool is a web-based application designed to size different workloads such as VDI, VSI, Microsoft SQL database and determine the appropriate Cisco HyperFlex systems. HX Tools are a highly

recommended suite of profiling, sizing, and configuration generating tools. For complete information on this tool and sizing guidelines, refer to the [HyperFlex Sizer Tool Site](#).

Considerations when Configuring Your Order

While configuring your order in CCW or related ordering tools, please take the following into consideration:

Considerations for new clusters:

- Use the HX Sizing Tools for generating an optimal and validated configuration
- HX Clusters need to use the same node types, CPU, and drives configurations
- When using SED Encrypted Drives, all drives in the node and cluster must be SED drives
- When using the HX Acceleration Engine, all nodes in the cluster must contain an HX Acceleration Engine

Refer to the Latest Release Notes document for the feature compatibility table and Software licenses required

Considerations for node or cluster expansions:

- HyperFlex is a scale out architecture so compute and storage are easily expanded.
- When adding drives to your nodes, use the same drive PID or the recommended drive PID from the appropriate table in the [HX Drive Compatibility](#) document.
- Expand the cluster with the same generation of server nodes or later, match CPU subsystem and confirm compatible drives from the appropriate table in the [HX Drive Compatibility](#) document.

For additional information, refer to the [HX Drive Compatibility](#) document.

Cisco HyperFlex Systems M5 Offerings


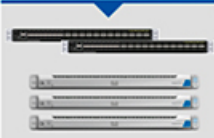



Cisco HyperFlex M5 Node Configuration Options

You can order HyperFlex nodes in 4 configurations:

- **Cisco HyperFlex Edge nodes for remote and branch offices**—orderable as HX-E-M5S-HXDP, and available in All Flash (HXAF-E-220M5SX -Cisco HyperFlex All Flash Edge 220 M5 system) or Hybrid (HX-E-220M5SX Cisco HyperFlex Hybrid Edge 220 M5 system).
- **Cisco HyperFlex Hybrid nodes**—HX220c M5 and HX240c M5 available in 1RU c220 (UCS M5 servers), 2RU 240 with Small Form Factor (SFF), and Hybrid LFF—2RU 240 with Large Form Factor (LFF) drives (highest storage capacity system). LFF is a subset of Hybrid. Both have spinning disk hard disk drive storage. Hybrid nodes use SSD drives for caching and HDDs for the capacity layer.
- **Cisco HyperFlex All-Flash nodes (HXAF)**—available in 220 SFF (HXAF220c All-Flash M5) and 240 SFF (HXAF240c All-Flash M5). All Flash nodes use SSD drives or NVMe storage for caching, and SSD drives for the capacity layer.
- **Cisco HyperFlex All-NVMe nodes**—orderable within two bundles: HXAF2X0C-M5S (with HX-FI-6332 or HX-FI-6454 Fabric Interconnects) or HXAF-M5S-HXDP (nodes without FI) and available as HXAF220C-M5SN Cisco HXAF220c M5 All NVMe HyperFlex System. This is the highest performance system. All NVME nodes use Intel Optane cache drives with NVME capacity drives.

The following illustration describes the cluster configuration options available for the Cisco HyperFlex Hybrid M5 nodes.

Figure 1: Cisco HyperFlex Hybrid M5 Configurations

HX220c M5 Edge Cluster	HX220c M5 Cluster	HX240c M5 Cluster	HX240c M5 LFF Cluster	HX M5 + Compute Node Clusters
				
*6.7TiB - 18.0TiB	**6.0TiB - 128.5TiB	**6.0TiB - 369.5TiB	**30.1TiB - 214.2TiB	NOTE: Consult Release Notes for Compute Node Support Details
Smallest Footprint 3 Node Cluster (VSI, ROBO)	Smallest Footprint 3-32 Node Cluster (VDI, ROBO)	Capacity-Heavy 3-32 Node Cluster (VDI & VSI Workloads)	Capacity-Heavy 3-8 Node Cluster (high capacity Workloads)	Compute-Heavy Hybrid (Compute Bound Apps/VDI)
Per-Node 1 x Cache SSD 3-8 x 1.8TB Capacity HDDs	Per-Node 1 x Cache SSD 6-8 x 1.2TB or 1.8TB Capacity HDDs SED Options Available	Per-Node 1 x Cache SSD 6-23 x 1.2TB or 1.8TB Capacity HDDs Support up to 2 GPUs SED Options Available	Per-Node 1 x Cache SSD 6-12 x 6TB or 8TB Capacity HDDs Support up to 2 GPUs	3-32 HX220 or HX240 Node Cluster + Up to 32 Compute Nodes Blade or Rack Local Disk, SD Card or SAN Boot

*Edge does not support RF3. Usable capacity w/ RF2 before compression and deduplication

**Usable capacity w/ RF3 before compression and deduplication

The following table describes the cluster configuration options for Cisco HyperFlex All-Flash and All NVMe M5 nodes.

Table 2: Cisco HyperFlex All Flash Edge, All Flash 220 Converged, All Flash 240 Converged, and All NVMe M5 Configurations

	HXAF220c M5 Edge Cluster	HXAF220c M5 Cluster	HXAF240c M5 Cluster	All NVMe
Nodes per Cluster	2,3,4 node cluster (VSI, ROBO)	3-32 node cluster (ESXi)	3-32 node cluster (ESXi)	3-32 node cluster (ESXi)
Max Stretched Cluster Size	Not Supported	Min: 2 x2, Max: 16 x 16 (HXDP-P Enterprise license required)	Min: 2 x2, Max: 16 x 16 (HXDP-P Enterprise license required)	Min: 2 x2, Max: 16 x 16 (HXDP-P Enterprise license required)
Max Capacity Drives per Node	Per Node: 3-8 x 960GB or 3.8TB or 7.6TB Capacity SSD	Per Node: 6-8 x 960GB or 3.8TB or 7.6TB Capacity SSD	Per Node: 6-23 x 960GB or 3.8TB or 6-12 x 7.6TB (max 12 drives per node)	1x Optane CACHE 6-8 1TB, 4TB, or 8TB capacity drives
Cluster Size Storage Capacity	5.76TB - 243TB	17.3TB - 1.95PB	17.3TB - 2.9PB	18TB - 2.05PB

1 RU formfactor nodes based on UCS C220 Rack Servers

Cisco HyperFlex HX220c M5 Hybrid, HXAF220c M5 All-Flash Nodes, and HXAF220-M5SN All NVMe Nodes—these nodes provide up to 8 capacity drives in a 2-socket, 1-rack unit (1RU) package ideal for small-footprint environments.

Cisco HyperFlex HX220c M5 Edge Nodes—these nodes are designed to work in simplified clusters using existing Gigabit Ethernet networks without Fabric Interconnects. Cisco HyperFlex Edge systems are easily deployed and managed with Intersight and also offer the same easy deployment and management as all Cisco HyperFlex systems.

2 RU formfactor nodes based on UCS C240 Rack Servers

Cisco HyperFlex HX240c M5 Hybrid and HXAF240c M5 All-Flash Nodes—these provide high disk capacity (up to 23 capacity drives) in a 2-socket, 2RU package ideal for storage-intensive applications.

Local Storage

The following table lists categories of local storage and how they are utilized by HyperFlex nodes.

Type of Drive	Description
Capacity Drive (Data)	<p>All NVMe nodes use NVMe drives and are the highest performance and lowest latency configuration.</p> <p>All Flash (HXAF) nodes use solid state drives (SSD), which do not have any moving parts, therefore enabling higher performance and reliability.</p> <p>Hybrid (HX) nodes use spinning hard disk drives (HDD) for data storage that provide the lowest cost per Gb of storage.</p> <p>HyperFlex also supports self-encrypting drives (SED) on HXAF and HX nodes. When ordering SED, ensure that the entire HyperFlex cluster utilizes SED encryption.</p> <p>Note HyperFlex does not support mixing SED and non-SED drives in the same cluster.</p>
Cache Drive (Write Log)	<p>Each HX node requires one high-performance drive for data caching and rapid acknowledgment of write requests.</p> <p>All NVMe nodes use Optane 3D Crosspoint drives for caching.</p> <p>All Flash (HXAF) nodes may use SSD or NVMe drives for caching.</p> <p>Hybrid (HX) nodes use SSD drives for caching.</p>
System Drive (Housekeeping)	HX and HXAF nodes use one System Log drive per system.

Drive Support

The following table describes drive support per product.

		PID	# per system	System type					
				HX220C-M5X	HXAF220C-M5X	HX240C-M5X	HXAF240C-M5X	HX240C-M5L	HXAF240C-M5L
Drive Controller			1 per system						
	M5	HX-SAS-M5		✓	✓			✓	

		PID	# per system	System type					
				H22CMBK	H22DMBK	H22CMBK	H22DMBK	H22C-MEL	H22DMB
	M5	HXSASM5HD				✓	✓		
DATA Capacity			Varies per system type	min 6 / max 8	min 6 / max 8	min 6 / max 23	min 6 / max 23	min 6 / max 12	
	M5	HFD21BKLN		✓		✓			
	M5	HFD81BKLN		✓		✓			
	M5	HFD21BKLN		✓		✓			
	M5	HFD21NKN9		✓		✓			
	M5	HFD67K4KN						✓	
	M5	HFD87K4KN						✓	
	M5	HFD27K4KN						✓	
	M5	HSD816XEV			✓		✓		
	M5	HSD906XEV			✓		✓		
	M5	HSD81BEN9			✓		✓		
	M5	HSD81BIN9			✓		✓		
	M5	HSD90BEN9			✓		✓		
	M5	HSD90BIN9			✓		✓		
	M5	HANME2100						✓	✓
	M5	HANME2100						✓	✓
	M5	HANMH1400						✓	✓
	M5	HSD7616XEV							
	M5	HSD61EMN9							
Cache			1 per system						
	M5	HSD1612XFP			✓	✓	✓		
	M5	HSD161BEN9				✓			
	M5	HSD90BEN9		✓	✓		✓		

		PID	# per system	System type					
				HX20C-M5X	HX20D-M5X	HX20E-M5X	HX20F-M5X	HX20C-M5L	HX20D-M5N
	M5	HX-M2-12XFP						✓	
	M5	HX-M2-12XFP			✓		✓		
	M5	HX-M2-12XFP			✓		✓		✓
System drive or aka Hosted drive			1 per system						
	M5	HX-M2-240GB		✓	✓	✓	✓	✓	
	M5	HX-M2-12XFP							✓
Boot Device	M5	HX-M2-240GB	1 per system	✓	✓	✓	✓	✓	✓
Internal microSD card	M5	HX-MSD-32G	1 per system	✓	✓	✓	✓	✓	✓



Note Power Cord Compatibility for India

New replacement parts for cords in compliance with the Dec 1st, 2020 Bureau of Indian Standards (BIS) regulatory standard for lower-power consumption products have been added.

For a complete list of part numbers, see the [Cisco HyperFlex Specification Sheets](#).

Cisco HyperFlex Cluster Sizing Guidelines

The Cisco HyperFlex Sizer tool is a web-based application designed to size different workloads such as VDI, VSI, Microsoft SQL database and determine the appropriate Cisco HyperFlex systems. HX Tools are a highly recommended suite of profiling, sizing, and configuration generating tools. For complete information on this tool and sizing guidelines, refer to the [HyperFlex Sizer Tool Site](#).

Considerations when Configuring Your Order

While configuring your order in CCW or related ordering tools, please take the following into consideration:

Considerations for new clusters:

- Use the HX Sizing Tools for generating an optimal and validated configuration
- HX Clusters need to use the same node types, CPU, and drives configurations
- When using SED Encrypted Drives, all drives in the node and cluster must be SED drives

- When using the HX Acceleration Engine, all nodes in the cluster must contain an HX Acceleration Engine

Refer to the Latest Release Notes document for the feature compatibility table and Software licenses required

Considerations for node or cluster expansions:

- HyperFlex is a scale out architecture so compute and storage are easily expanded.
- When adding drives to your nodes, use the same drive PID or the recommended drive PID from the appropriate table in the [HX Drive Compatibility](#) document.
- Expand the cluster with the same generation of server nodes or later, match CPU subsystem and confirm compatible drives from the appropriate table in the [HX Drive Compatibility](#) document.

For additional information, refer to the [HX Drive Compatibility](#) document.

