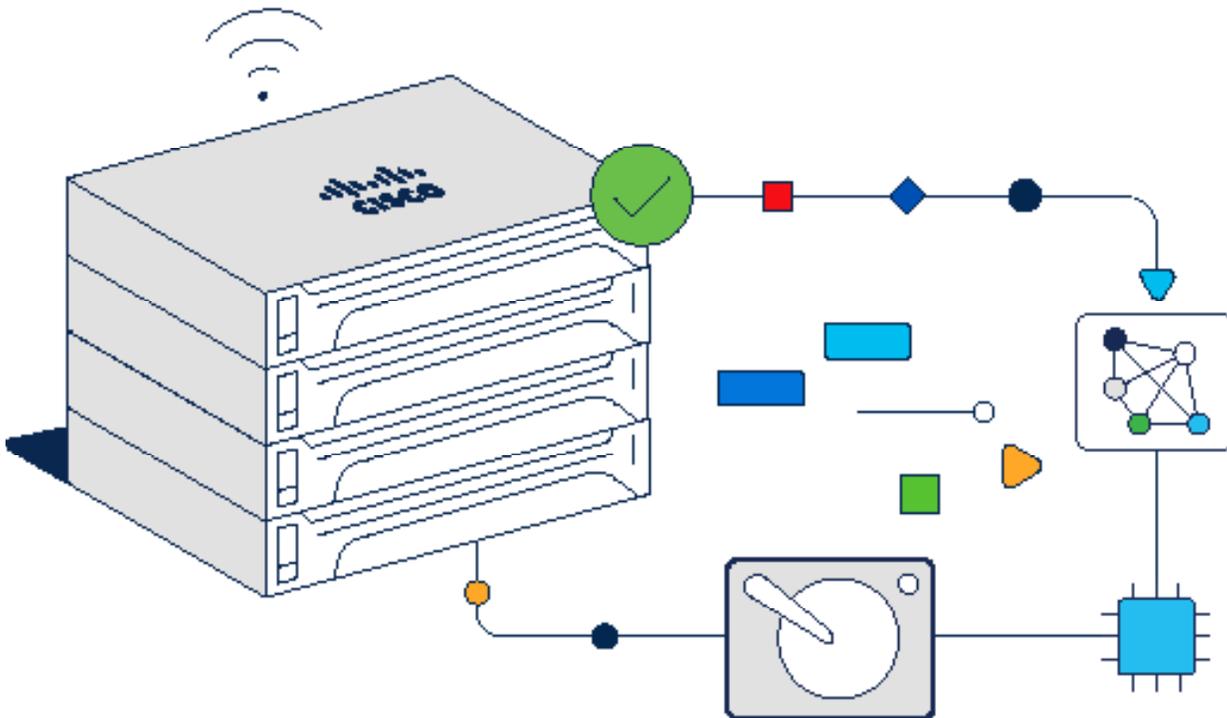


## Cisco HyperFlex

# Modernize your digital transformation with HyperFlex

Cisco HyperFlex unlocks the full potential of hyperconvergence for a wide range of workloads and use cases. Cisco HyperFlex is fast to deploy, simple to manage, easy to scale, and ready to provide a unified pool of resources to power all your business applications.



# Contents

<b>Introduction .....</b>	<b>3</b>
Cisco HyperFlex Overview .....	4
HyperFlex Capabilities .....	5
Hypervisor support .....	5
Kubernetes integration .....	5
Data protection .....	6
Data services .....	7
Thin provisioning .....	7
Snapshots .....	8
Native replication .....	8
Stretch clusters .....	8
Fast, space-efficient clones .....	8
Enterprise-class availability .....	9
Data rebalancing .....	9
Network platform .....	9
Operations .....	10
Edge .....	11
Security .....	12
Hardware .....	12
<b>Conclusion .....</b>	<b>13</b>
<b>Research from Gartner .....</b>	<b>13</b>



## Introduction

Digital transformation is disrupting every industry by breaking down barriers between people, business, and things. By breaking these barriers, they can create new products and services and find more efficient ways of doing business. IT is at the heart of every company, but its role has evolved from helping a company run more efficiently to ultimately helping it grow faster. Increasing customer demands, the exponential growth in the number of devices, a mobile workforce, and multicloud-delivered applications are making infrastructure harder to manage. To address these complexities and remain relevant, IT must digitize the services it offers. Line-of-business leaders are looking to transform their business processes and achieve tangible outcomes. However, to do this, they need help from IT leaders who can provide them with technology solutions to achieve these outcomes. As we navigate these uncertain times, almost all industries are innovating across the core, edge, and cloud to deal with the rapid change of technology and a more dispersed workforce to deliver services and an enriching experience to customers and their workforce.

Current IT infrastructure weighs data centers down with sluggish arrays of hard-to-manage components, preventing them from keeping up with the quick-turnaround demands of business and the cloud. By centralizing resources and management, hyperconverged infrastructure (HCI) lowers costs, reduces complexity, and alleviates staff burdens while improving performance. In essence, HCI is designed to work as one system with software-managed storage, as opposed to the separate components of converged solutions.

HCI simplifies administration by providing a single point of management. It fully integrates with your entire data center, eliminating the need for separate servers and network storage and delivering on-demand infrastructure for data-centric workloads.

IT is continually tasked to do more with fewer resources. In addition to taking care of existing IT infrastructure, it must drive new initiatives such as AI, machine learning (ML), and big data. For large and small organizations, hyperconvergence's centralized management, scalable architecture, and cloudlike simplicity deliver multiple benefits.



Choosing an HCI platform means assessing its technical capabilities in great depth and determining whether it is well suited to your unique mix of workloads. Accordingly, Gartner has developed a list of evaluation criteria to inform and help direct your technical decision making. This summary highlights Cisco HyperFlex capabilities against Gartner's key criteria.



## Cisco HyperFlex Overview

Cisco HyperFlex combines Cisco UCS networking and computing technology, and the Cisco HyperFlex Data Platform to deliver a complete, preintegrated solution. Cisco HyperFlex supports all-flash-memory, all-NVMe, and hybrid-storage configurations and provides PCIe slots capable of supporting GPUs. In most hyperconverged environments, networking is an afterthought. HyperFlex delivers complete hyperconvergence with networking as an integral and essential part of its system. Using Cisco UCS fabric interconnects, you have a single point of connectivity and management that incorporates HX-series nodes and Cisco UCS servers: a key architectural element that no other hyperconverged vendor can offer. After you deploy a cluster, you can scale it to its maximum size without needing to redesign the network. The solution is designed for easy, smooth scaling. Hyperconverged systems need massive amounts of east-west traffic bandwidth and low latency, and Cisco UCS fabric interconnects deliver both with up to 100 Gbps networking.

Cisco HyperFlex offers a great deal of choice in implementation such as:

- Compute only - You can add converged or compute-only nodes to expand an HX cluster.
- Fabric interconnects - a choice of networking architectures and topologies that can include FI or No-FI as needed.
- Edge - Specific nodes designed for edge infrastructure that brings the robust feature set of Cisco HyperFlex systems to edge environments with high-capacity storage and cloud management in a small footprint while leveraging existing networking gears.
- HX Express - Our most popular node configurations with simple and important options, priced attractively to deliver optimal value, and reduced transaction times to help keep plans on track.
- SW Only - We have announced that we will be introducing a software-only version of HyperFlex that provides customers the flexibility to run Hyperflex software on existing and new Cisco UCS servers, third-party and white box servers that include ruggedized form factors, and eventually in the public cloud, enabling a consistent experience across the data center, edge, and cloud.
- Cisco+ - Provides all the benefits of Cisco, now as a service with the hybrid cloud offer giving you flexible consumption for your on-premises infrastructure so you can optimize workloads with the flexibility to choose an OpEx consumption model tailored to the needs of the business - pay as you use or pay as you grow.



## HyperFlex Capabilities



### Hypervisor support

#### Gartner criteria:

Every HCI solution supports at least one hypervisor, and many support more than one. However, what set of hypervisors a particular HCI supports varies widely among different products and vendors. When evaluating HCIS offerings, therefore, it is important to consider the range of hypervisors that a product supports and its level of integration with each. Gartner evaluates hypervisors based on compute efficiency, virtual machine (VM) performance, VM resilience, and operations.

#### Cisco HyperFlex capabilities:

HyperFlex is foundationally designed to be multi-hypervisor ready and currently supports VMware vSphere ESXi and Microsoft Hyper-V hypervisors. HyperFlex supports ESXi and Hyper-V-native capabilities for high availability, data protection, and business continuity. vSphere vMotion and Windows Hyper-V Live Migration are supported, while I/O striping and HyperFlex's distributed file system ensures minimal performance impact of migrated VMs.



### Kubernetes integration

#### Gartner criteria:

Containers are the other major form factor in which software-defined compute resources can be delivered. They have assumed a critical role in application delivery and portability. All the major HCI stacks include container management as an integral part of the platform. However, the need for container management is broader than just the ability to create and destroy a container. State-of-the-art infrastructure architectures rely on orchestrated fleets of containers, and Kubernetes is the container orchestrator of choice. In the leading HCI platforms, Kubernetes serves as a means of ensuring application portability between public and private infrastructure and of managing internal clusters. The leading HCI products provide their own turnkey Kubernetes implementations and also have deep integrations with Kubernetes management constructs for networks and storage.

#### Cisco HyperFlex capabilities:

HyperFlex supports any container management framework that uses Kubernetes as container orchestration including Redhat OpenShift Docker Datacenter, Rancher, PKS, etc., running in virtualized environments. Using programmable APIs and single sign-on, HyperFlex management integrates with container technologies to enable cloud-native applications running in microservices architecture.



## Data protection

### Gartner criteria:

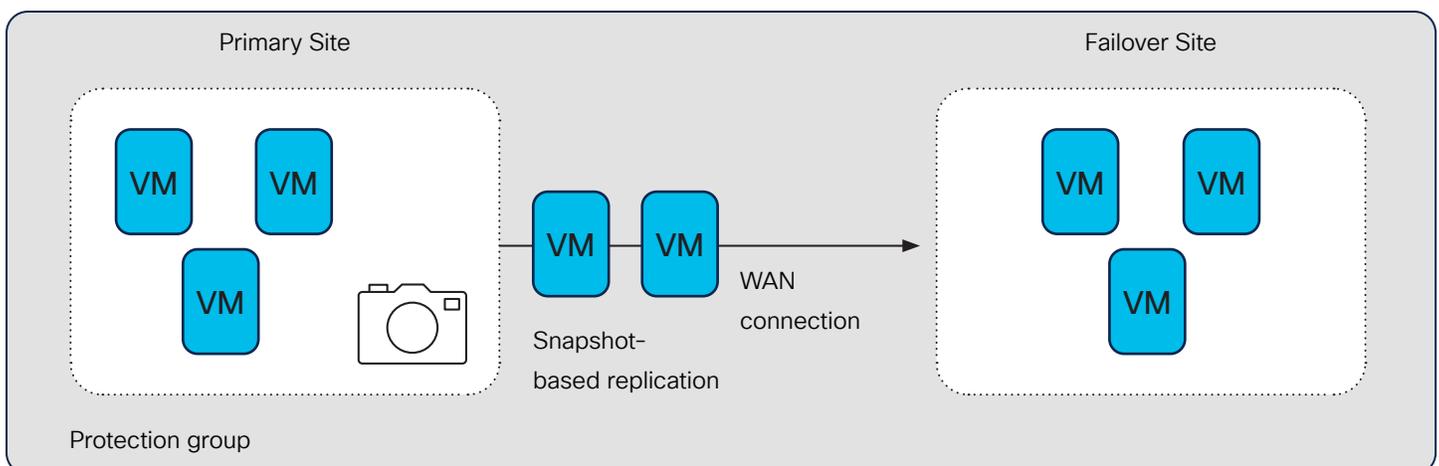
Data protection is the most important goal of any storage device. Data loss in production is never acceptable. Therefore, to be considered enterprise-ready, any storage device must take steps to make data redundant and highly available. Every HCI vendor takes data protection seriously, and every HCI platform has some HA mechanism for data; however, there are vast differences between products, and many different HA-enabling technologies are at play.

### Cisco HyperFlex capabilities:

Cisco HyperFlex-native replication can create a point-in-time copy of your business-critical application data at a remote site without the need for matching infrastructure at both sites, thus saving cost. Your remote site can be a smaller, simpler version of your primary site. In fact, you can use all-flash storage at your primary site and hybrid storage at your remote site. This approach reduces complexity as well as costs while protecting your data.

With Cisco HyperFlex-native replication, you can protect a single virtual machine or sets of virtual machines through protection groups using Cisco HyperFlex Connect. To define the data to be replicated, you set the policy in the protection groups to manage the same way that periodic snapshots are performed. You can choose the interval at which replications occur, which can range from 15 minutes to 24 hours (see figure below). Replication can also be bidirectional, so you can protect source virtual machines at the destination and destination virtual machines at the source. This capability is especially useful if your remote location is also an active production site.

In addition, Cisco HyperFlex systems partners with leading data protection vendors (such as Veeam, Cohesity, and Commvault) to work together to protect data. Running this software on Cisco HyperFlex systems provides additional recovery functions, delivering application-consistent snapshots and replication between hosts.



Rather than streaming data through a single control node, Cisco HyperFlex–native replication streams data from each node at the primary site to all the nodes at the recovery site for fast and efficient data transfer. This approach to moving data also helps prevent hot spots at the replication site because the replication workload is automatically distributed evenly across all the data nodes. In addition, the data is compressed and moves in larger blocks to move data more efficiently to your secondary site. The solution also offers configurable bandwidth management features, enabling you to control the volume of data traversing the wire so that it does not overwhelm your network.

Lastly, a HyperFlex stretched cluster is a single cluster with nodes in two geographically distributed sites. Both sites of the cluster act as primary for certain user VMs. The data for these VMs is replicated synchronously on the other site. Stretched clusters enable you to access the entire cluster even if one of the sites were to completely go down. Typically, these sites are connected with a low latency, dedicated, high-speed link between them. HyperFlex Stretched Cluster enables you to deploy an Active–Active disaster avoidance solution for mission critical workloads requiring high uptime (near zero Recovery Time Objective) and no data loss (zero Recovery Point Objective).



## Data services

### Gartner criteria:

Any enterprise–grade storage platform provides a rich set of services for the data it houses, and HCI storage platforms are no exception. However, there is perhaps more variability between HCI offerings in this category

than in any other. Assessing each product’s set of data services often reveals important and fundamental differences between two competing HCI products. Data services are often the most important differentiators in head-to-head comparisons of HCI stacks, and customers often make their purchasing decisions based on them

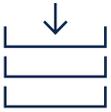
### Cisco HyperFlex capabilities:

The Cisco HyperFlex Data Platform provides a scalable implementation of space–efficient data services, including thin provisioning, pointer-based snapshots, native replication, and clones–without affecting performance.



## Thin provisioning

The platform makes efficient use of storage by eliminating the need to forecast, purchase, and install disk capacity that may remain unused for a long time. Virtual data containers can present any amount of logical space to applications, whereas the amount of physical storage space that is needed is determined by the data that is written. As a result, you can expand storage on existing nodes and expand your cluster by adding more storage–intensive nodes as your business requirements dictate, eliminating the need to purchase large amounts of storage before you need it. Lastly, the Cisco HyperFlex HX Data Platform makes your data instantly accessible and highly available, with always–on deduplication and compression that reduces your storage needs.



## Snapshots

The Cisco HyperFlex Data Platform uses metadata-based, zero-copy snapshots to facilitate backup operations and remote replication: critical capabilities in enterprises that require always-on data availability. Space-efficient snapshots allow you to perform frequent online backups of data without needing to worry about the consumption of physical storage capacity. Data can be moved offline or restored from these snapshots instantaneously.



## Native replication

This feature is designed to provide policy-based remote replication for disaster recovery and virtual machine migration purposes. You add virtual machines to protection groups that inherit the policies you define. Native replication can be used for planned data movement (for example, migrating applications between locations) or unplanned events such as data center failures.

Unlike enterprise-shared storage systems, which replicate entire volumes, we replicate data on a per-virtual-machine basis. This way you can configure replication on a fine-grained basis so that you have remote copies of the data you care about.

The data platform coordinates the movement of data with the remote data platform, and all nodes participate in the data movement using a many-to-many connectivity model. This model distributes the workload across all nodes, avoiding hot spots and

minimizing performance impacts. Once the first data is replicated, subsequent replication is based on data blocks changed since the last transfer.



## Stretch clusters

With stretch clusters you can have two identical configurations in two locations acting as a single cluster. With synchronous replication between sites, a complete data center failure can occur, and your applications can still be available with zero data loss. In other words, applications can continue running with no loss of data. The recovery time objective is only the time that it takes to recognize the failure and put a failover into effect.



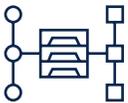
## Fast, space-efficient clones

Clones are writable snapshots that can be used to rapidly provision items such as virtual desktops and applications for test and development environments. These fast, space-efficient clones rapidly replicate storage volumes so that virtual machines can be replicated through just metadata operations, with actual data copying performed only for write operations. With this approach, hundreds of clones can be created and deleted in minutes. Compared to full-copy methods, this approach can save a significant amount of time, increase IT agility, and improve IT productivity.



## Enterprise-class availability

In the Cisco HyperFlex HX Data Platform, the log-structured distributed-object layer replicates incoming data, improving data availability. Based on the policies you set, data that is written to the write cache is synchronously replicated to one or more caches located in different nodes before the write operation is acknowledged to the application. This approach allows incoming write operations to be acknowledged quickly while protecting data from storage device or node failures. If an SSD, NVME device, or node fails, the replica is quickly recreated on other storage devices or nodes using the available copies of the data.



## Data rebalancing

A distributed file system requires a robust data-rebalancing capability. In the Cisco HyperFlex Data Platform, no overhead is associated with metadata access, and rebalancing is extremely efficient. Rebalancing is a nondisruptive online process that occurs in both the caching and capacity layers, and data is moved at a fine level of specificity to improve the use of storage capacity. The platform automatically rebalances existing data when nodes and drives are added or removed or when they fail. When a new node is added to the cluster, its capacity and performance are made available to new and existing data. The rebalancing engine distributes existing data to the new node and helps ensure that all nodes in the cluster are

used uniformly from both capacity and performance perspectives. If a node fails or is removed from the cluster, the rebalancing engine rebuilds and distributes copies of the data from the failed or removed node to available nodes in the clusters.



## Network platform

### Gartner criteria:

HCI network platforms permit highly abstracted programmatic control of network infrastructure. A principal goal of many of these platforms is easing the management burden of hybrid and multicloud networking. For example, many HCI platforms can build overlay networks that provide a common IP address space across on-premises and public cloud infrastructure. Some HCI stacks even manage network underlays, with direct control of physical switches and other network devices such as load balancers or firewalls.

### Cisco HyperFlex capabilities:

Networking is important in hyperconverged systems because the data platform's performance depends on it. With Cisco UCS fabric interconnects, you get high-bandwidth, low-latency, unified fabric connectivity that carries all production IP, hyperconvergence-layer, and management traffic over a single set of cables. Every connection in the cluster is treated as its own virtual link, with the same level of security as if it were supported with a separate physical link. This makes the

integrated network more secure than when commodity approaches are used. Additionally, topologies support single top-of-rack (ToR) and dual ToR switch options for ultimate network flexibility and redundancy.

The system is designed so that all traffic, even from blade server compute-only nodes, reaches any other node in the cluster with only a single network hop. This single-hop architecture accelerates east-west traffic, enhancing cluster performance. Our latency is deterministic, so you get consistent network performance for the data platform, and you don't have to worry about network constraints on workload placement.

Lastly, Cisco HyperFlex supports Cisco ACI and VXLAN integration. VXLAN is qualified, as a standalone protocol or in conjunction with ACI, for use with Stretched Cluster.



## Operations

### Gartner criteria:

The operational benefits of hyperconvergence have been a selling point for HCI since the very beginning of the technology. Nowadays, HCI vendors are more likely to tout their hybrid-cloud and multicloud management tools. Many HCI platforms can manage public cloud resources from the same management GUI as on-premises infrastructure. And HCI software now routinely includes full-featured cloud management platforms that provide provisioning, orchestration, inventory, classification, monitoring, and analytics for both public and private infrastructure. By definition, all HCI software includes a single, unified management

console for compute, networking, and storage. But HCI products take vastly different approaches to their management GUIs. Some have written proprietary, web-based consoles from the ground up. Others provide plug-ins for existing management tools such as vCenter. Equally important is programmatic access to the system.

### Cisco HyperFlex capabilities:

HyperFlex leverages Cisco Intersight to provide fully remote lifecycle management of multiple clusters right from installation, including upgrades, expansion, maintenance, and troubleshooting (including the network, compute, and storage pieces) from the cloud and from on-premises infrastructure. Intersight is architected as a capabilities-as-a-service platform sharing common backend services that provide single-sign-on (SSO), role-based access control (RBAC), a common inventory, and shared APIs and SDKs.

Intersight's cloud-native, microservice, SaaS architecture, and its dynamic CI/CD development model, allow us to innovate rapidly and deliver at an unprecedented pace - weekly updates automatically roll out for all our customers. This has some industry-changing advantages:

- Deployment and upgrade at scale - Customers can now deploy and upgrade their HyperFlex clusters from the cloud, in parallel, through lights-out management functionality and using policies to ease configuration. These policies can be reused to quickly scale from a single cluster to thousands. Customers can choose to do this either through the UI or automate it using the REST APIs.

- Cluster Expansion – With Intersight SaaS based management for HyperFlex, customers can remotely expand their HyperFlex clusters with converged or compute only nodes. This enables customers to start as small as a two nodes edge cluster and expand exponentially as the business grows.
- CI/CD pipeline – Intersight rolls out new capabilities every two weeks on average. Customers using Intersight have the latest and most up-to-date software, significantly reducing calls to Cisco Technical Assistance Center (TAC) and eliminating “known issues.”
- Dashboards and alerts – These enable customers to monitor their global operating environment at a glance from any location.
- Launch point – From the cluster list view, customers can dig deeper into individual clusters as needed. This enables management from a laptop, tablet, or mobile device from anywhere in the world.
- HCL check – Intersight evaluates the compatibility of HyperFlex systems to check if the hardware and software have been tested and validated by Cisco or Cisco partners. Intersight reports validation issues after checking the compatibility of the server model, processor, firmware, adapters, operating system, and drivers and displays the compliance status with the hardware compatibility list (HCL).
- Connected TAC – A portfolio of services for IT, developers, and application teams to minimize time to innovation with: Intersight Kubernetes Service, Intersight Workload Optimizer, Intersight Infrastructure Services, Intersight Service for HashiCorp Terraform, and many more..



## Edge

### Gartner criteria:

Hyperconvergence has always been an enabling technology for edge services. The ability to support two-node clusters and permit the witness node to reside in the cloud.

### Cisco HyperFlex capabilities:

HyperFlex Edge requires a minimum of two nodes. This ensures that high availability and fault tolerance are maintained without any single point of failure. In this configuration, either node can be taken down for service proactively for hardware and software upgrades or for troubleshooting. This configuration also allows for unexpected failures, because the cluster is designed to continue to run on a single node without disrupting access to the storage system, data services, and running workloads. HyperFlex Edge upgrades are orchestrated through Intersight and provide seamless full-stack upgrade capability across multiple sites in parallel. These online upgrades take advantage of the fact that the two-node cluster can tolerate a node failure and continue servicing all workloads.

A key design tenant for HyperFlex Edge includes the use of Intersight for end-to-end lifecycle management across a geographically dispersed HCI infrastructure. Only two-node HyperFlex Edge clusters require an external witness, and this capability is provided by the Cisco Intersight Invisible Cloud Witness feature. A first in the HCI space, the cloud witness is provided

as a SaaS offering hosted by Cisco. Customers can choose a Cisco cloud-hosted or on-premises version of Intersight to support remote witnessing. Key benefits include elimination of a third site and infrastructure to host a witness VM, cloud-like simplicity with zero configuration and auto-upgrades, secure communications using well known ports, and a custom-built silent witness protocol – enabling almost any WAN to be used with the cloud witness.



## Security

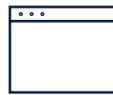
### Gartner criteria:

Security is a critical – and nonoptional – consideration in enterprises of any size. To be enterprise-ready, HCI systems must ensure that the data they store remains secure and that the system itself remains uncompromised. Because any software-defined technology separates the control and data planes, security for each plane must be addressed individually and the vulnerabilities of each treated separately.

### Cisco HyperFlex capabilities:

Cisco HyperFlex uses a policy-based approach to security to ensure uniform, consistent, compliant, and secure encryption management and deployment

across a cluster. In order to support enterprise-critical applications, HyperFlex deploys a holistic approach that integrates security deep into the platform. Data-at-rest encryption complies with regulations that require the use of security best practices. HyperFlex also provides a hardened platform based on a secure development lifecycle that protects against vulnerabilities and threats.



## Hardware

### Gartner criteria:

Hyperconverged products are often delivered as appliances – that is, servers prepackaged with HCI software. However, hyperconvergence is a software solution, and even in an HCI appliance, the principal value of a system is the software it runs. And HCI software requires physical entities on which to run: racks of servers, plus network switches connecting them. The hardware should also support NVMe drives, GPUs, and OpEx-based financing models.



## Conclusion

Every organization will have unique requirements and priorities. Cisco HyperFlex Systems deliver a complete, next-generation hyperconverged solution for a wide range of applications and in any location, including your enterprise data center, remote edge locations, private cloud, and hybrid cloud environments.

Cisco HyperFlex systems unlock the full potential of hyperconvergence so that you can use a common platform to support more of your applications and

use cases, including virtualized and containerized applications, virtual desktops, big data environments, enterprise applications, databases, and internet infrastructure applications, along with test and development environments. Cisco HyperFlex systems deliver the operational requirements for agility, scalability, and OpEx consumption models – but with the benefits of on-premises infrastructure.

## Research from Gartner

Digital transformation is pushing enterprises to be more distributed and to expand to the edge. Information and operations leaders must plan for the growth of IoT and related data at the edge, protect their enterprise from edge growing pains in security and connectivity, and prepare for changes in edge computing use cases.

[Predicts 2022: The Distributed Enterprise Drives Computing to the Edge](#)

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

C96-3151008-00 11/22

