

# Cisco HyperFlex Systems

## Manage an explosion of applications

Today's applications are more diverse and distributed than ever, living across a complex, distributed, multidomain world that encompasses enterprise data centers; campus, branch, and edge locations; and private and public clouds. Because these applications form your organization's persona, you need to support the explosion in the number and types of applications, enable cloud-native approaches, and help to drive an ever greater digital reach.

Cisco HyperFlex™ systems with Intel® Xeon® Scalable processors deliver hyperconvergence with power and simplicity for any application, anywhere. Engineered with Cisco Unified Computing System™ (Cisco UCS®) technology, and managed through the Cisco Intersight™ cloud operations platform, Cisco HyperFlex systems can power your applications and data anywhere, optimize operations from your core data center to the edge and into public clouds, and increase agility by accelerating DevOps practices. We accomplish this through:

- **An application-centric hyperconverged platform** that delivers any application to any location, and at any scale, predictably and securely
- **A cloud operating model** that enables you to manage and optimize your environment—on premises, edge, and public cloud—from anywhere, with the power of data, analytics, and cloud-based management
- **Adaptable infrastructure** that adjusts to your ever-changing business needs, accelerating innovation by enabling your developers while helping to preserve security and compliance for the business

Cisco HyperFlex is certified as an Intel® Select Solution for Hybrid Cloud.



Cisco HyperFlex systems with Intel Xeon Scalable processors



## Contents

<b>The challenge: A digitally driven economy</b>	<b>3</b>
<b>The solution: Cisco HyperFlex systems</b>	<b>3</b>
Bridge the gap .....	4
What makes Cisco HyperFlex systems different .....	5
Complete end-to-end solution .....	6
Ready for all of your applications .....	6
Engineered on Cisco UCS technology .....	7
Cisco Intersight cloud operations platform.....	7
Powered by next-generation data technology .....	8
Agile, efficient, adaptable, and scalable solution.....	8
<b>Solution architecture</b>	<b>9</b>
Cluster composition with Cisco Hyperflex HX-Series nodes .	9
Complete hyperconvergence with Cisco networking .....	10
Integration with Cisco Application Centric Infrastructure.....	11
<b>Cisco HyperFlex HX Data Platform</b>	<b>11</b>
Cisco HyperFlex HX Data Platform controller .....	13
Data distribution.....	14
Logical availability zones .....	15
Data read and write operations .....	15
Enterprise-class storage features .....	16
<b>Engineered on Cisco UCS technology</b>	<b>17</b>
Cisco HyperFlex HX-Series Nodes .....	17
Independent resource scaling with Cisco UCS servers.....	18
Hyperconvergence at the network edge .....	18
Connect to external shared storage.....	20
Powered by Intel Xeon Scalable processors .....	20
Manage your HyperFlex clusters.....	20
<b>Conclusion</b>	<b>21</b>
Part of our broad data center strategy .....	21

## What you will learn

This document is designed to help explore the benefits and architecture of Cisco HyperFlex systems:

- [“The solution: Cisco HyperFlex systems” on page 3](#) provides a six-page overview, with an emphasis on how our hyperconverged system is a more agile, efficient, adaptable, and scalable solution.

If the overview leaves you wanting to learn more, read [“Solution architecture” on page 9](#), with topics including:

- [“Cisco HyperFlex HX Data Platform” on page 11](#) provides an overview of the platform’s clustered, distributed file system, including the enterprise storage features it supports, how it integrates with the hypervisor, how data is distributed, and how the various levels of caching contribute to high performance.
- [“Engineered on Cisco UCS technology” on page 17](#) discusses the benefits of using Cisco UCS technology as the foundation for Cisco HyperFlex, including flexible node configurations, independent resource scaling, our edge offering that can be deployed at scale, and how the platform’s integrated networking forms the backbone of the cluster including connectivity to enterprise shared storage.

## The challenge: A digitally driven economy

The shift toward a digitally driven economy has accelerated the number and type of applications that organizations are deploying. With applications increasingly representing the face of business, these trends cannot be overlooked. In its [FutureScape: Worldwide IT Industry 2020 Predictions](#), IDC projects that by 2023:

- **300 percent more applications** will run in the data center and edge locations.
- **500 million digital applications** and services will be developed using cloud-native approaches.
- **More than 40 percent** of new enterprise IT infrastructure will be deployed at the edge.

The increasing diversity and distribution of applications poses a new challenge to IT organizations. With 46 percent of enterprise products and services expected to be digital, or digitally delivered, by 2022 ([IDC, 2019](#)), your business will rely ever more on IT infrastructure. You need solutions that help bridge the gap between IT and business demands with the capability to:

- **Power apps and data anywhere.** Containers are the virtual machines of this decade, so you must be ready to support cloud-native software with cloud-like resource delivery. You need to deploy these applications with any scale and reach, and you still need to host enterprise applications and virtual desktop environments running in traditional virtual machines.
- **Optimize operations.** Today’s cloud-native applications deploy literally anywhere, in your core data center, at the edge, and in one or more public clouds. You need tools that help you deploy into hybrid cloud environments with application-aware analytics that help maintain performance levels through intelligent resource management.
- **Increase agility.** You need to be able to deliver the right resources to your applications wherever they reside. This takes open, future-proof infrastructure that enables rapid iteration through DevOps processes while preserving security and maintaining compliance for the business.

## The solution: Cisco HyperFlex systems

To help you meet the challenge of deploying large numbers of new applications at any scale, and in any location, we developed Cisco HyperFlex™ systems. In today’s world this adaptable platform acts as your on-premises and edge infrastructure that complements and integrates with the workloads that you deploy into public clouds. Tight integration with the Cisco Intersight™ cloud operations platform enables full lifecycle management of your workloads wherever you want to deploy them, locally, at the edge, and into the cloud. With management hosted in the cloud, you have access to unlimited deployment locations and scale.

## Meet evolving challenges

### Simplify the core

“Over the last two years, IT organizations spent 70% on ‘run the business’ IT spending, up from 67% in 2013 to 2014, and 65% in 2012.”

04/18 Gartner, 2018 Strategic roadmap for compute infrastructure, Chirag Dekate, Daniel Bowers, et. al., April 10, 2018.

### Deploy cloud-native applications

“By 2022, 90% of new enterprise applications will be developed as cloud-native applications developed with agile methodologies and based on a hyperagile API-based architecture that leverages microservices architectures, containers, and serverless functions.”

04/18 Gartner, 2018 Strategic roadmap for compute infrastructure, Chirag Dekate, Daniel Bowers, et. al., April 10, 2018.

### Reach to the edge

“By 2022, more than 50% of enterprise-generated data will be created and processed outside [of] the core data center or cloud.”

12/18 Gartner, Gartner IT Infrastructure, Operations & Cloud Strategies Conference Event Presentation, The Future of IT Infrastructures: Always On, Always Available, Everywhere, David J. Cappuccio, Bob Gill, et. al., December 3–6, 2018.

## Bridge the gap

Our strategy is to help you bridge the gap by providing the IT capabilities you need in order to help your business thrive in an always-on world (Figure 1):

- **App-centric platform:** We help you deliver any app, to any location, at any scale, both predictably and securely. This starts with infrastructure that delivers cloud-like resource delivery that complements what you get from the cloud. This helps you differentiate your services from the competition.
- **Cloud operations platform:** A cloud operating model helps you manage distributed operations at scale, from physical and virtual infrastructure deployment to workload placement and resource optimization based on real-time analysis of your application performance. This delivers true IT as a service to your business, helping support the drive to deliver more applications in more locations.
- **Adaptable infrastructure:** We support open, future-proof infrastructure you to support your applications. A hyperconverged application platform optimized to deliver cloud-native apps delivered as microservices Traditional application hosting with both VMware vSphere and Microsoft Windows Server Hyper-V virtual machines. This approach supports the DevOps processes that your organization is embracing and opens the door to more growth opportunities.

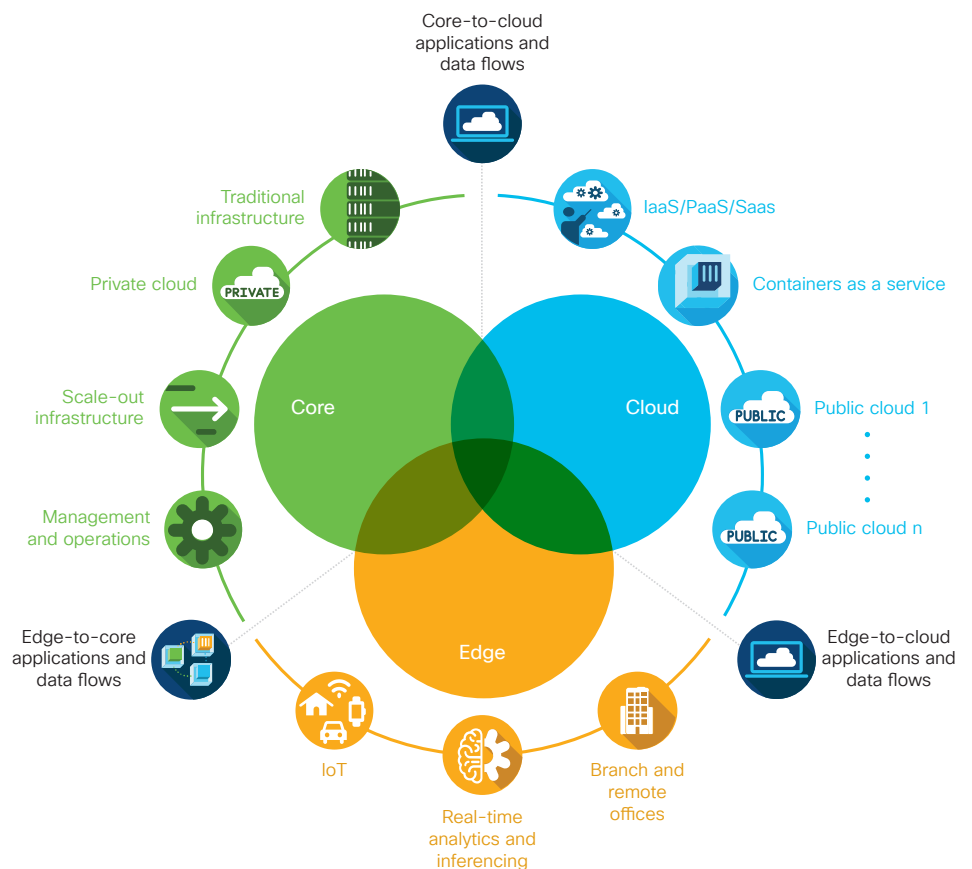


Figure 1 Cisco HyperFlex systems support the data center core, cloud, and edge.

# Certified Intel Select Solution

Cisco HyperFlex is now certified as an Intel® Select Solution for Hybrid Cloud. Built on Intel® Xeon® Scalable processors and using Intel® Optane™ Solid State Drives, HyperFlex has been verified as optimized infrastructure for hybrid cloud strategies.

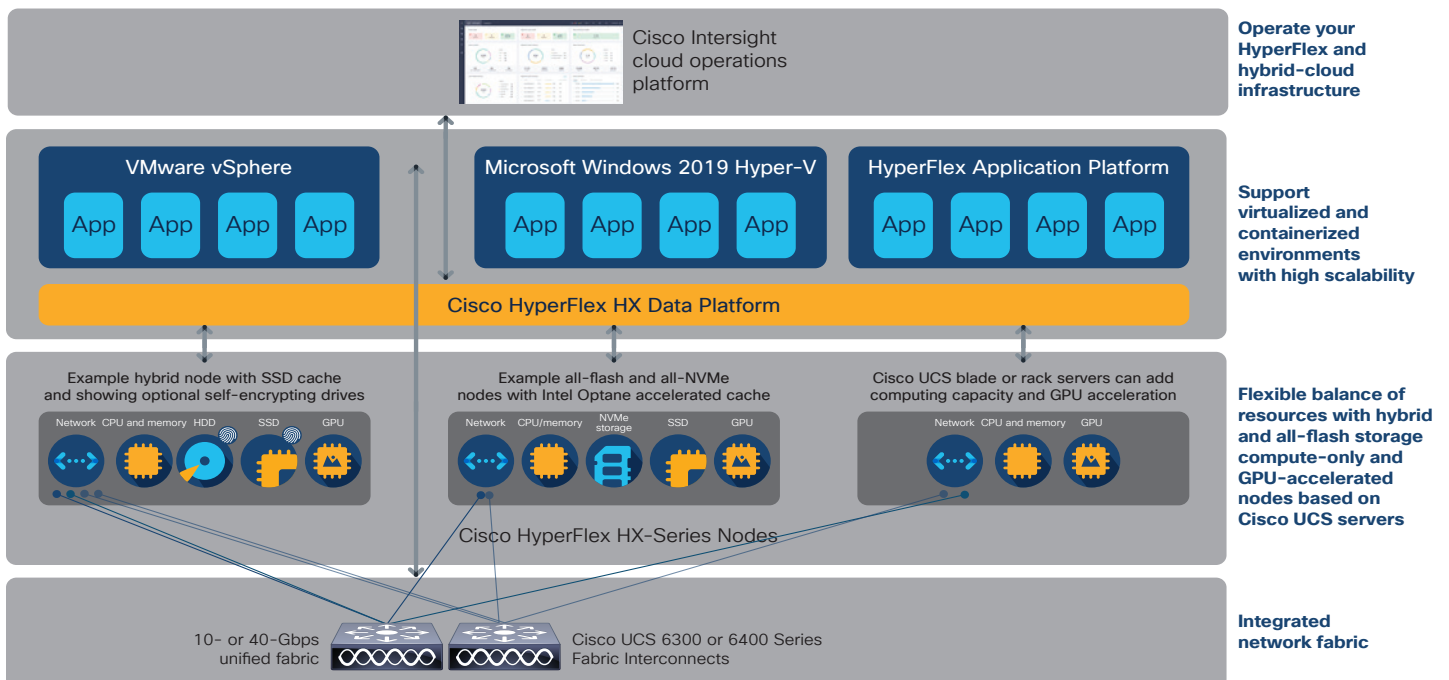


## What makes Cisco HyperFlex systems different

Cisco HyperFlex unlocks the full potential of hyperconvergence for a wider range of workloads and use cases. Our innovations include:

- **Enterprise-grade application platform** with a fully integrated, multitenant Kubernetes environment with full lifecycle management provided by the Cisco Intersight cloud operations platform.
- **Fully integrated all-NVMe nodes** that are engineered for performance and reliability
- **A unique hardware data compression accelerator**
- **Full bare-metal deployment and software maintenance** managed from the cloud
- **Flexible GPU acceleration** in virtual desktop environments
- **Powerful GPU acceleration** for AI and ML inferencing applications.

The Cisco HyperFlex platform is faster to deploy, simpler to manage, easier to scale, and ready to provide a unified pool of resources to power your business applications. Cisco HyperFlex systems integrate into the data center you have today to form the foundation of your hybrid-cloud strategy. You can deploy Cisco HyperFlex systems wherever you need them, from central data center environments to remote locations and edge-computing environments (Figure 2).



**Figure 2** Cisco HyperFlex systems support virtualized and containerized applications and let you choose the exact combination of resources to power your enterprise applications

## What's new?

Cisco HyperFlex systems support:

- **Cisco HyperFlex Application Platform.** The software provides an enterprise-grade, multitenant, 100 percent upstream Kubernetes environment without the cost and complexity of licensing underlying virtualization software and manually integrating hardware and software components. Because the platform uses the integrated data platform, the complexity of persistent storage for containers is eliminated. The underlying storage is provided by the Cisco HyperFlex HX Data Platform that protects your data with all of the platform's enterprise storage features (see ["Cisco HyperFlex HX Data Platform" on page 11](#))
- **Cisco Intersight Kubernetes Service.** The Cisco Intersight platform provides full lifecycle management of your application platform with rapid deployment, orchestration, management, and monitoring of containerized environments. These features enable IT organizations to strike a balance between giving developers the power to specify their own containerized infrastructure and the needs of IT organizations to uphold service-level agreements with the reliability and availability needed to maintain them.

## Complete end-to-end solution

We engineered Cisco HyperFlex systems to support a broader range of applications and workloads in the data center, private and hybrid clouds, remote locations, and edge-computing environments. This new generation extends the ease of hyperconverged system deployment, management, and support beyond your central data center to hybrid cloud environments and to the network edge.

Cisco HyperFlex systems combine:

- **Software-defined computing** in the form of nodes based on Cisco Unified Computing System (Cisco UCS) servers
- **Software-defined storage** provided by the powerful Cisco HyperFlex HX Data Platform software
- **Software-defined networking** with Cisco Unified Fabric that integrates smoothly with Cisco® Application Centric Infrastructure (Cisco ACI™)
- **Cloud-based management** with Cisco Intersight that supports full lifecycle management—from initial deployment to workload placement and optimization to single-click, full-stack software upgrades.

Together, these elements comprise an adaptive infrastructure that lets you integrate easily with your existing infrastructure. The result is a cluster that comes up and configures itself quickly and that scales resources independently to closely match your application resource needs.

## Ready for all of your applications

The power and performance of Cisco HyperFlex brings lower latency and increased readiness to support a wider range of applications than ever before. We have invested in developing Cisco Validated Designs that have tested and validated many traditional enterprise applications in the Cisco HyperFlex environment to help you speed deployment and reduce risk. These guidebooks for implementation help you accelerate deployment and reduce risk for virtual desktop environments (Citrix or VMware), Epic Electronic Health Record (EHR) environments, Oracle Database, Microsoft SQL Server, big data applications including Splunk and SAP HANA, and graphics-accelerated high-performance computing, artificial intelligence (AI), and machine learning (ML) applications. With flexible configurations, you can deploy Cisco HyperFlex systems in your enterprise data center, to create private clouds, or for edge computing.

The rapid shift to cloud-native development models means that you must establish an on-premises complement to public-cloud deployments so that you can maintain security and compliance with workloads and data that must reside locally. The Cisco Intersight Kubernetes Service provides a ready-to-consume Kubernetes environment that speeds and simplifies containerized applications.

## What's new, continued

- **Cisco Intersight Workload Optimizer.** This feature helps ensure continuous health of applications across on-premises and cloud environments. The intelligent software continuously analyzes workload consumption, costs, and compliance constraints and automatically allocates resources in real time to support application performance. You can determine when, where, and how to move and resize workloads, maximize elasticity with public cloud resources, and quickly model infrastructure and workload growth scenarios to determine how much infrastructure you will need and when you will need it.
- **Core data center enhancements.** The data platform now supports iSCSI storage, with initiators in Microsoft Windows, Linux, virtual machines, physical servers, and container persistent volumes able to draw on this block-based storage approach.
- **Edge computing enhancements.** New, more storage- and GPU-dense enable support for more data at the edge, and N:1 native backup helps keep your data safe in your core data center as well. Security enhancements further harden edge locations.

### Engineered on Cisco UCS technology

Cisco UCS provides a single point of connectivity and hardware management that integrates Cisco HyperFlex nodes into a single unified cluster. The system is self-aware and self-integrating so that when a new component attached, it is automatically incorporated into the cluster. Rather than requiring you to configure each element in the system manually through a variety of element managers, Every aspect of a node's personality, configuration, and connectivity is set through management software.

You can independently scale the computing power of your cluster by configuring Cisco UCS servers as computing-only nodes or compute- and GPU-only nodes. This feature lets you adjust the balance of CPU, GPU acceleration, and storage capacity to tune cluster performance to your workloads without the additional costs associated with additional nodes. Incremental scalability allows you to start small and scale as your needs grow. You can accelerate virtual desktop environments with graphics processing units (GPUs) that integrate directly into nodes, speeding graphics-intensive workloads for a smooth user experience.

### Cisco Intersight cloud operations platform

Cisco Intersight provides an essential control point for organizations. With Intersight, you can get more value from hybrid IT by simplifying operations across on-premises data centers, edge sites, and public clouds, continuously optimizing and accelerating service delivery.

Cisco Intersight is a unified, secure, software-as-a-service platform comprised of modular services that bridge applications with infrastructure. It provides correlated visibility and management across bare-metal servers, hypervisors, Kubernetes, and serverless and application components, helping you transform your operations with artificial intelligence to reach the scale and velocity your stakeholders demand.

Intersight helps your teams securely collaborate and work smarter and faster together, within and across domains, by automating lifecycle workflows and enabling consistency and governance with extensible, open capabilities that natively integrate with third-party platforms and tools. With Intersight, your teams can intelligently visualize, optimize, and orchestrate all of your applications and infrastructure anywhere they are.

In the context of Cisco HyperFlex systems, Intersight gives you instant access to all of your clusters regardless of where they are deployed. Parallel, heterogeneous deployment supports the scale you need to support distributed locations, and full-stack upgrade allows you to update the entire stack—firmware, hypervisor, and data platform software—to the revision levels your applications need. High-level resource inventory and status are provided by Intersight dashboards. Intersight Kubernetes Service configures and manages 100 percent upstream Kubernetes environment supported by the HyperFlex Application Platform. The Intersight Workload Optimizer monitors

## More agile

Cisco HyperFlex systems are more agile because they perform, scale, and interoperate better:

- **Deployment is fast and easy.** Your cluster ships with the software preinstalled and ready to launch. Installation is the same whether your cluster is in the core data center or at a remote site on the network edge.
- **Integrated networking brings high performance.** Your cluster is interconnected with low, consistent latency, and with 10- and 40-Gbps network bandwidth.
- **Scaling is fast and simple.** The system automatically discovers new hardware when it is installed. Then adding it to the cluster takes only a few mouse clicks. Logical availability zones allow clusters to grow to up to 64 nodes while minimizing the impact of multiple node failures.
- **Interoperability is straightforward.** Management capabilities enable you to install and operate your Cisco HyperFlex system in the data center you have today, with high-level management tools that support operations across your hyperconverged and your traditional infrastructure.

application workloads and helps to optimize resource allocation to maintain quality-of-service levels. A recommendation engine helps you proactively respond to impending issues such as the need to scale capacity. Connected Cisco TAC integration can automatically open a support case when errors are detected and can even upload log information needed for case analysis. The Intersight invisible witness supports lightweight 2-node clusters and helps maintain continuous operation in the event of a node or network failure with fully automatic configuration. Storage analytic capabilities can track and monitor compute, network, and storage configurations and provide proactive optimization recommendations.

### Powered by next-generation data technology

The Cisco HyperFlex HX Data Platform combines the cluster's storage devices into a single distributed, multitier, object-based data store. It makes this data available through the file system protocols and mechanisms needed by the higher-level hypervisors, virtual machines, and containers.

Performance scales linearly as you scale the cluster because all components contribute both processing and storage capacity to the cluster. The data platform optimizes storage tiers for an excellent balance between price and performance. For example, hybrid nodes use solid-state disk (SSD) drives for caching and hard-disk drives (HDDs) for capacity; all-flash nodes use fast SSD drives or Nonvolatile Memory Express (NVMe) storage for caching and SSDs for capacity; all-NVMe nodes deliver the highest performance for the most demanding workloads with caching further accelerated by Intel Optane™ SSDs.

A full set of enterprise-class data management features are built in: for example, snapshots, thin provisioning, replication for disaster recovery and backup, data encryption, integration with third-party backup tools, and instant and space-efficient cloning. Cisco UCS fabric interconnects secure your data through its lifecycle, with security and compliance controls for protection when you distribute, migrate, and replicate data across storage environments. Security features help you comply with industry and governmental standards. The platform delivers high availability through parallel data distribution and replication, accelerated by the low latency and high bandwidth of the Cisco Unified Fabric. Data is continuously optimized with deduplication, compression, and optional encryption, helping reduce your storage costs without affecting performance. Dynamic data placement in server memory, caching, and capacity tiers increases application performance and reduces performance bottlenecks.

### Agile, efficient, adaptable, and scalable solution

Bringing benefits to your IT organization and to your business, Cisco HyperFlex systems are agile, efficient, and adaptable, making them well suited for hosting environments such as virtual desktops, server virtualization deployments, and test and development environments.

## More efficient

Our solution is more efficient because of the following:

- **Containers without the cost.** The HyperFlex Application Platform supports an enterprise-grade Kubernetes cluster without the expense of licensing hypervisor software.
- **Buy only the storage you need.** Continuous data deduplication and compression, fast, space-efficient clones, thin provisioning, and optional hardware-accelerated compression all contribute to lowering the cost of your storage.
- **Data protection is built in.** You can use native HX Data Platform snapshot and replication capabilities with the same data protection solutions you use in your data center because Cisco HyperFlex systems interoperate with leading backup tools.
- **Your data is secure.** You don't need to take extra steps to secure your data at rest when using self-encrypting drives. Cisco HyperFlex systems help you save time and ensure compliance.
- **Save on storage administration.** You don't need to install a complex storage network or worry about logical unit numbers (LUNs). Enterprise shared storage can connect directly to your cluster if needed.

- **More agile.** Cisco HyperFlex systems are more agile because they perform, scale, and interoperate better (see [“More agile” on page 8](#))
- **More efficient.** Our solution was designed from the beginning with a purpose-built, highly efficient data platform that combines the cluster's scale-out storage resources into a single, distributed, multitier, object-based data store. Features you expect of enterprise storage systems are built into Cisco HyperFlex systems (see [“More efficient” on page 9](#)).
- **More adaptable.** Your business needs and your workloads are constantly changing. Your infrastructure needs to quickly adapt to support your workloads and your business (see [“More adaptable” on page 11](#)).
- **More scalable.** Clusters can scale to up to 64 nodes, with protection from multiple node and component failures (see [“More scalable” on page 12](#)).

## Solution architecture

Cisco HyperFlex systems combine Cisco UCS networking and computing technology, powerful Intel Xeon Scalable processors, and the HX Data Platform to deliver a complete, preintegrated solution. After you install locally or through the Intersight interface, your cluster is ready to work for you whether you need to support virtualized or containerized applications. You get a uniform pool of computing, networking, and storage resources that is designed to power your applications. When you need to provision computing or storage capacity, that capacity is drawn from the entire pool.

### Cluster composition with Cisco Hyperflex HX-Series nodes

Physically, a cluster is composed from a set of three or more HX-Series nodes. These are integrated to form a single system through a pair of Cisco UCS 6300 or 6400 Series Fabric Interconnects with 10- or 40-Gbps connectivity (Figure 3). Cisco UCS servers can contribute additional computing and GPU acceleration capacity, participating in the data platform but contributing only computing power (and possibly graphics acceleration) to the cluster. These compute-only nodes can be configured up to a ratio of two compute-only nodes per Cisco HyperFlex node.

Cisco HyperFlex Edge solutions are deployed as fixed sets of two, three, or four edge-specific nodes that use Cisco or third-party Gigabit or 10 Gigabit Ethernet switches, offering the utmost in flexibility for deployment in remote and branch-office environments (see [“Hyperconvergence at the network edge” on page 18](#)).

A 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.

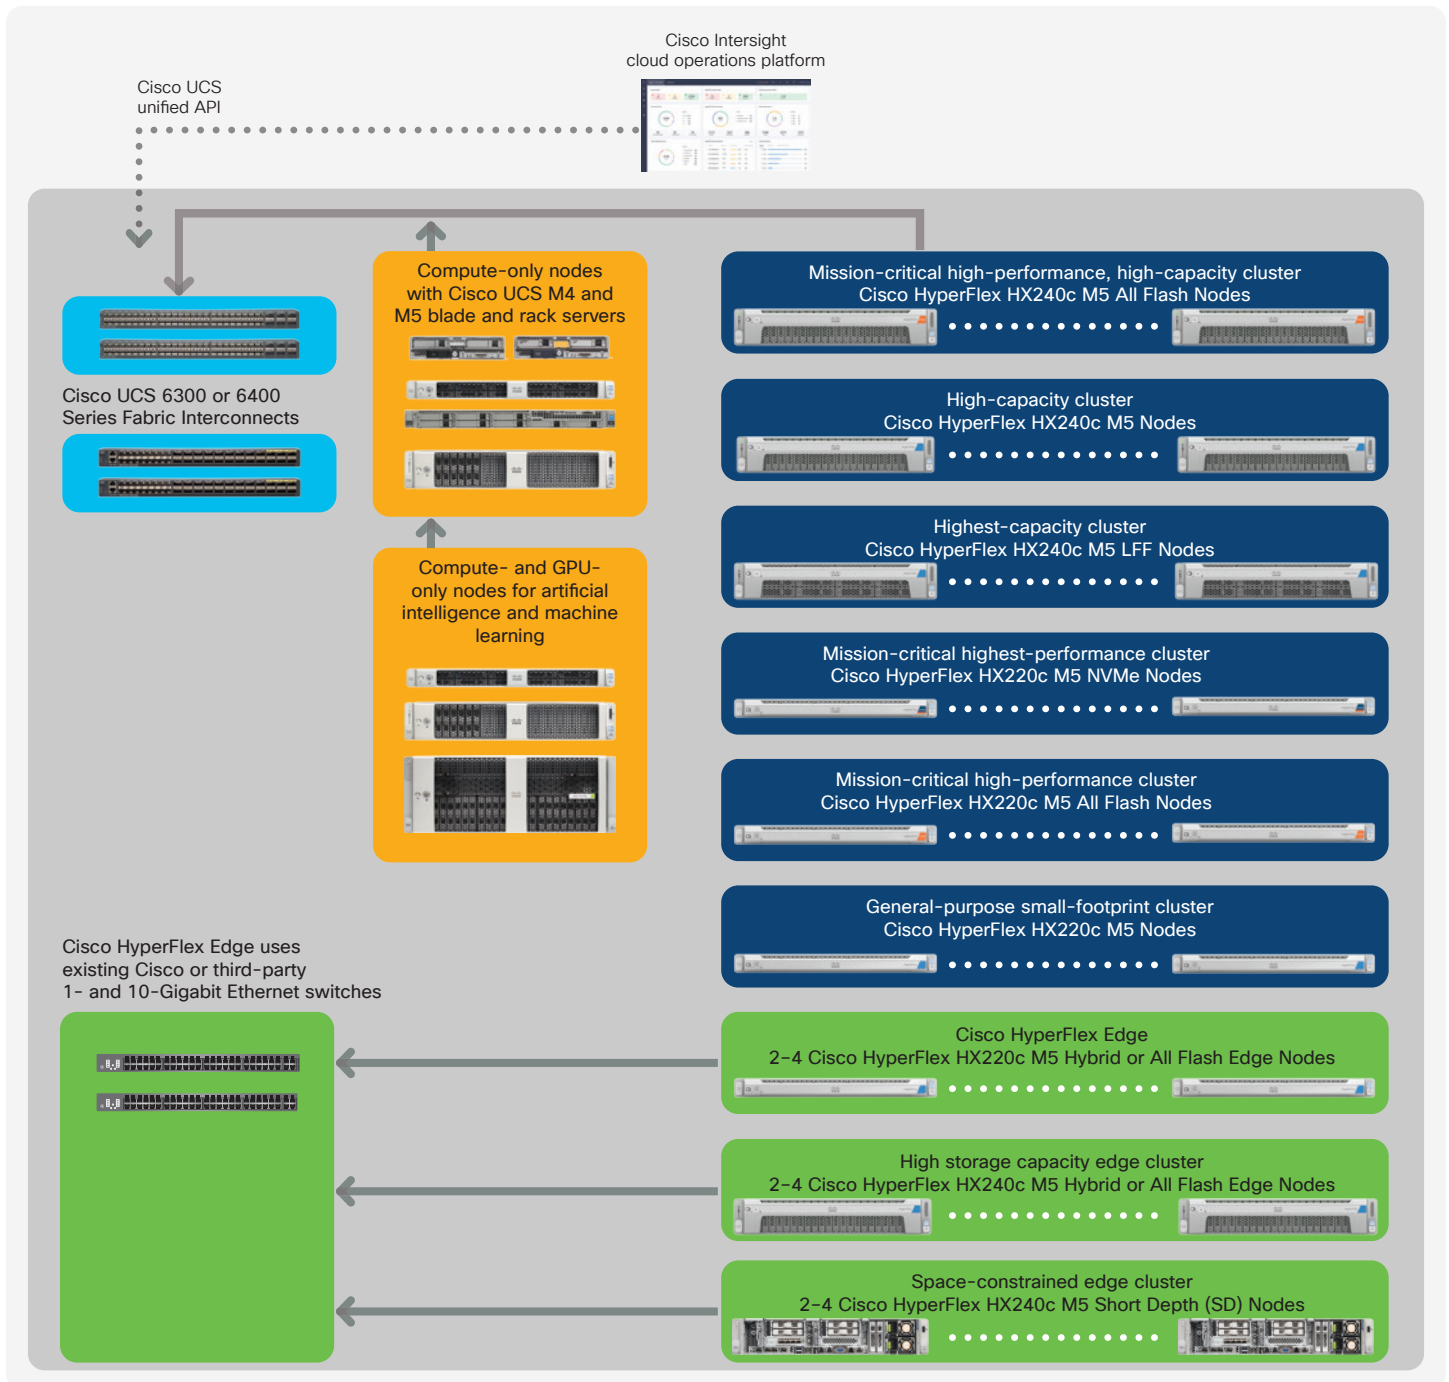


Figure 3 Cisco HyperFlex systems product family

### Complete hyperconvergence with Cisco networking

Networking in most hyperconverged environments is an afterthought. We give you complete hyperconvergence with networking as an integral and essential part of the 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

## More adaptable

Your business needs and your workloads are constantly changing. Your infrastructure needs to quickly adapt to support your workloads and your business.

- **Easy resource expansion and contraction.** You can scale resources up and out without having to adjust your software or networking capabilities.
- **Multisite support.** You can stretch a cluster across two sites for immediate recovery from disasters. Or you can distribute clusters geographically and replicate data between primary and backup sites.
- **Nondisruptive scaling.** Your infrastructure can easily scale out without the need to take down your cluster to add a node.
- **AI and ML support.** You can adapt your clusters to support AI and ML workloads by configuring up to 6 GPUs in Cisco HyperFlex HX240c Nodes, up to 2 GPUs in edge nodes, and GPU-only nodes to further augment the processing capability of your cluster.
- **Pay as you grow.** You can grow in small increments that won't break your budget, and you can independently scale your computing and capacity resources so they adapt to fit your specific application needs.

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 either 10- or 40-Gbps networking.

Networking is important in hyperconverged systems because the data platform 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.

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.

### Integration with Cisco Application Centric Infrastructure

As your environment grows and begins to span your enterprise, you can use Cisco ACI to implement a software-defined network. Cisco ACI provides automated, policy-based network deployment that secures your applications within secure, isolated containers. The network can attach directly to virtual machines and physical servers with increased security, real-time monitoring and telemetry, and automated performance optimization. You get consistency at scale because you can use Cisco ACI to interconnect your entire data center network, including your hyperconverged cluster.

## Cisco HyperFlex HX Data Platform

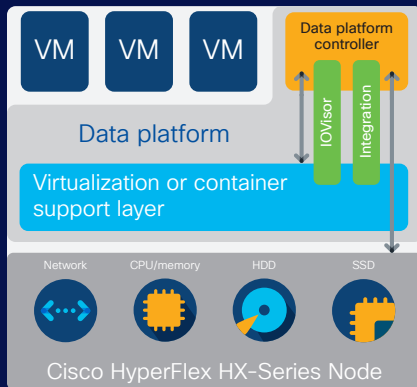
The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, scale-out file system with a wide array of enterprise-class data management services. The data platform's innovations redefine distributed storage technology, giving you complete hyperconvergence with enterprise storage features. The platform supports the leading hypervisors and also the Cisco HyperFlex Application Platform, bringing enterprise-class storage to cloud-native applications. The platform provides:

- **Enterprise-class data management features** required for complete lifecycle management and enhanced data protection in distributed storage environments. Features include deduplication, compression, thin provisioning, fast space-efficient clones, snapshots, native replication, and integration with backup solutions from leading vendors.

## More scalable

You need a hyperconverged solution that is up to the scale of today's enterprise and big data applications that need large numbers of nodes to support your business.

- **More nodes.** We have doubled the maximum scale to 64 nodes, with a maximum of 32 Cisco HyperFlex nodes and 32 Cisco UCS compute-only nodes. Stretch clusters can have 8 data and 8 compute-only nodes in each of two locations
- **More resiliency.** We have implemented logical availability zones that help you to scale without compromising availability.
- **More capacity.** You can choose nodes with large-form-factor disk drives for even greater capacity. This allows your cluster to scale to higher capacity for storage-intensive applications.
- **Multiprotocol support** includes providing direct iSCSI connectivity to the data platform. This provides enhanced support for Microsoft Windows Server failover clustering, Oracle software, Microsoft Exchange Server, and Kubernetes support (below). Support for iSCSI also extends access to external iSCSI devices to virtual machines supported on your HyperFlex cluster.
- **Integration with cloud-native applications** to supply persistent storage with enterprise-class features to enable stateful Kubernetes containers. Through an iSCSI-based Container Storage Interface (CSI) plugin, the platform delivers instant persistent volume creation with no need to move data when Kubernetes pods get rescheduled. This feature supports containers deployed with Intersight Kubernetes Service, native Kubernetes, Google Cloud's Anthos, and Red Hat OpenShift Container Platform.
- **Continuous data optimization** with always-on data deduplication and compression that does not affect performance, helping increase resource utilization with more headroom for data scaling.
- **End-to-end product security** includes developing the HyperFlex data platform software to be compliant with Cisco InfoSec standards which includes secure design, coding, and ongoing vulnerability analysis. The platform is secured from the firmware up with a secure root of trust. A reduced attack surface for the HyperFlex data platform helps to increase security of the entire platform.
- **Securely encrypted storage** optionally encrypts both the caching and persistent layers of the data platform. Integrated with enterprise key management software or with passphrase-protected keys, encryption of data at rest helps you comply with industry and government regulations. The platform itself is hardened to Federal Information Processing Standard (FIPS) 140-1, and the encrypted drives with key management comply with the FIPS 140-2 standard.
- **Dynamic data placement** optimizes performance and resilience by enabling all cluster resources to participate in I/O responsiveness. Hybrid nodes use a combination of SSD drives for caching and HDDs for the capacity layer; all-flash nodes use SSD drives or NVMe storage for the caching layer and SSDs for the capacity layer. NVMe nodes use even higher-performance storage.
- **Clusterwide parallel data distribution** implements data replication for high availability and performance, accelerated by the low latency and high bandwidth of the Cisco UCS network fabric.
- **Stretch clusters** span geographic distance with continuous operation in the event of a data center failure with no data loss.
- **Logical availability zones** increase scalability and availability because they automatically protect data against multiple component and node failures.
- **Native replication** creates remote copies of individual virtual machines for disaster recovery purposes. With all nodes in both the local and remote clusters participating in the replication activity, the overhead imposed on each cluster node is minimal. Replication for Edge clusters protect virtual



**Figure 4** Cisco HyperFlex Data Platform Controller plugs into the hypervisor in each node

machines from logical corruption or accidental deletion; it protects virtual machines from an edge cluster or site outage; and it supports migrating virtual machines from one edge site to another.

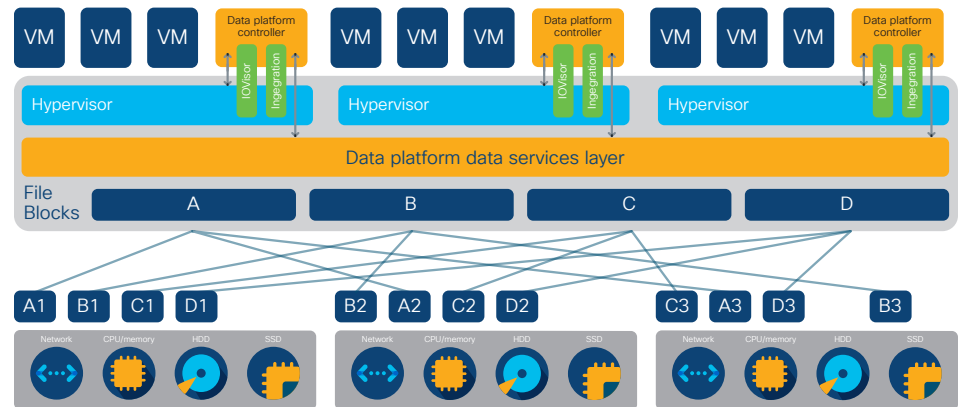
- **Linear and incremental scaling** provides faster, easier scalability compared to enterprise shared-storage systems, in which controller resources become a bottleneck and necessitate a complete upgrade of the storage system. Instead, whenever you add an increment of storage in Cisco HyperFlex systems, you also increment the data platform processing capacity.
- **API-based data platform architecture** provides data virtualization flexibility to support existing and new cloud-native data types. An API for data protection allows enterprise backup solutions to create snapshot-based backups of virtual machines.
- **A simplified approach** eliminates the need to configure LUNs or to require a storage administrator to configure SANs. Storage and data services can be managed through hypervisor tools such as VMware vCenter.

### Cisco HyperFlex HX Data Platform controller

A Cisco HyperFlex Data Platform controller resides on each node and implements a distributed scale-out file system (Figure 4). The controller runs in user space within a virtual machine and intercepts and handles all I/O from guest virtual machines. Dedicated CPU cores and memory allow the controller to deliver consistent performance without affecting performance of the other virtual machines in the cluster. When nodes are configured with self-encrypting drives, the controller negotiates with Cisco UCS Manager to receive the encryption keys that enable the drives to encrypt and decrypt data that flows to and from the various storage layers.

The data platform has modules to support the specific hypervisor or container platform in use. The controller accesses all of the node's disk storage through hypervisor bypass mechanisms for excellent performance. It uses the node's memory and dedicated SSD drives or NVMe storage as part of a distributed caching layer, and it uses the node's HDDs, SSD drives, or NVMe storage for distributed storage. The data platform controller interfaces with the hypervisor in two ways:

- **IOVisor:** The data platform controller intercepts all I/O requests and routes them to the nodes responsible for storing or retrieving the blocks. The IOVisor makes the existence of the hyperconvergence layer transparent to the hypervisor.
- **Hypervisor agent:** A module uses the hypervisor APIs to support advanced storage system operations such as snapshots and cloning. These are accessed through the hypervisor so that the hyperconvergence layer appears just as if it were enterprise shared storage. The controller accelerates operations by manipulating metadata rather than actual data copying, providing rapid response, and thus rapid deployment of new application environments.



**Figure 5** Data blocks are replicated across the cluster

### Data distribution

The HX Data Platform controller handles all read and write requests for volumes that the hypervisor accesses and thus intermediates all I/O from the virtual machines and containers. Recognizing the importance of data distribution, the HX Data Platform is designed to exploit low network latencies and parallelism, in contrast to other approaches that build on node-local affinity and can easily cause data hotspots.

With data distribution, the data platform stripes data evenly across all nodes, with the number of data replicas determined by the policies you set (Figure 5). This approach helps prevent both network and storage hot spots and makes I/O performance the same regardless of virtual machine location. This feature gives you more flexibility in workload placement and contrasts with other architectures in which a data locality approach does not fully utilize all available networking and I/O resources.

- **Data write operations:** For write operations, data is written to the local SSD or NVMe cache, and the replicas are written to remote caches in parallel before the write operation is acknowledged. Writes are later synchronously flushed to the capacity layer HDDs (for hybrid nodes) or SSD drives (for all-flash nodes) or NVMe storage (for NVMe nodes).
- **Data read operations:** For read operations in all-flash nodes, local and remote data is read directly from storage in the distributed capacity layer. For read operations in hybrid configurations, data that is local usually is read directly from the cache. This process allows the platform to use all solid-state storage for read operations, reducing bottlenecks and delivering excellent performance.

In addition, when migrating a virtual machine to a new location, the data platform does not require data movement because any virtual machine can read its data from any location. Thus, moving virtual machines has no performance impact or cost.

## Logical availability zones

As the maximum cluster size increases, the risk of a two-node failure shutting down the cluster increases. To mitigate this risk, logical availability zones partition the cluster into several logical zones. The HX Data Platform then ensures that there is no more than one copy of data in every zone. This feature can be enabled and automatically configured with the click of a mouse.

With logical availability zones, multiple node failures in a single zone don't cause the entire cluster to fail, because other zones contain data replicas.

## Data read and write operations

The data platform implements a log-structured file system that uses a caching layer to accelerate read requests and write responses, and it implements a capacity layer for persistent storage. The capacity layer is comprised of HDDs (in hybrid nodes), SSD drives (in all-flash nodes), or NVMe storage (in all-NVMe nodes).

Incoming data is striped across the number of nodes that you define to meet your data availability requirements. The log-structured file system assembles blocks to be written to a configurable cache until the buffer is full or workload conditions dictate that it be destaged to the capacity layer. When existing data is (logically) overwritten, the log-structured approach simply appends a new block and updates the metadata. When data is destaged, the write operation consists of a single seek operation with a large amount of data written. This approach improves performance significantly compared to the traditional read-modify-write model, which is characterized by numerous seek operations and small amounts of data written at a time.

When data is destaged to the capacity layer in each node, the data is deduplicated and compressed. This process occurs after the write operation is acknowledged, so there is no performance penalty for these operations. A small deduplication block size helps increase the deduplication rate. Compression further reduces the data footprint. Data is then moved to the capacity tier as cache segments become free (Figure 6).

Read operations in hybrid nodes cache data on the SSD drives and in main memory for high performance. In all-flash and NVMe nodes they read directly from storage. Having the most frequently used data stored in the caching layer helps make Cisco HyperFlex systems perform well for virtualized applications. When virtual machines modify data, the original block is likely read from the cache, so there is often no need to read and then expand the data on a spinning disk. The data platform decouples the caching tier from the capacity tier and allows independent scaling of I/O performance and storage capacity.

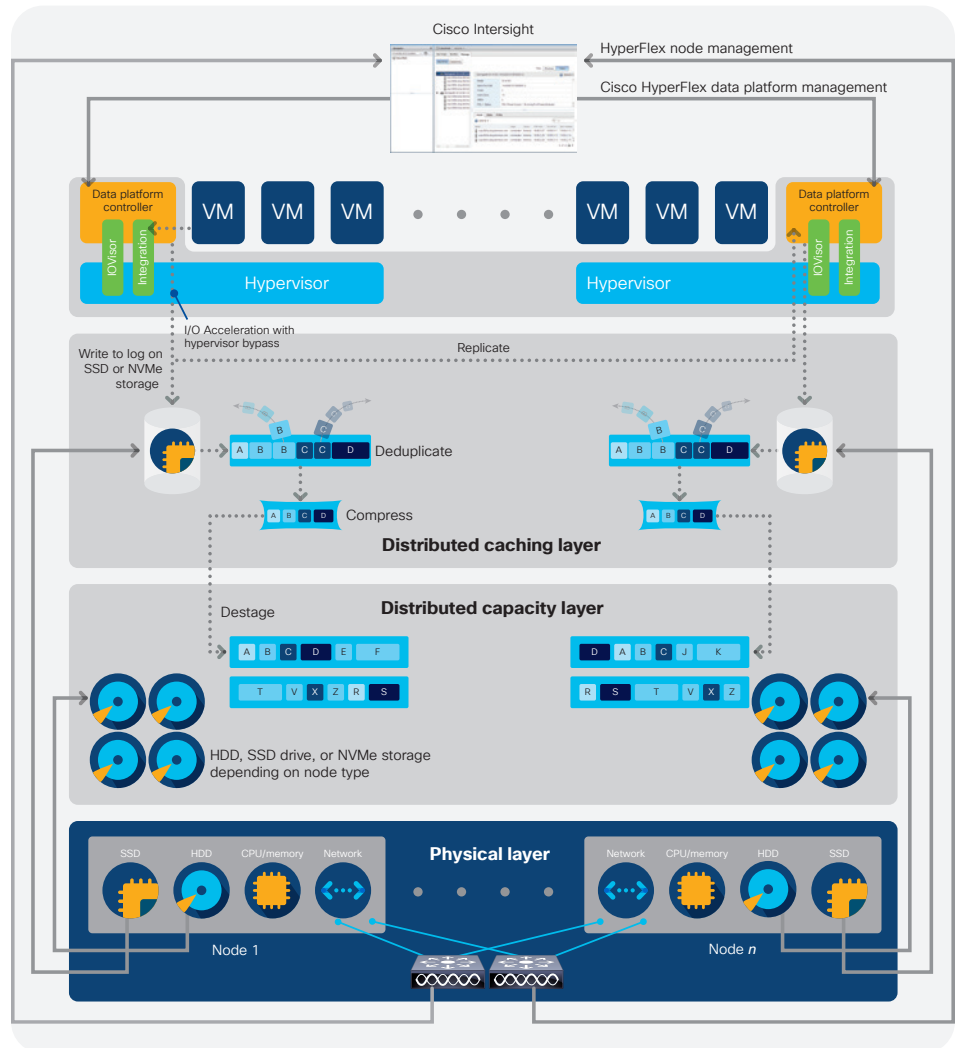


Figure 6 Data write flow through the Cisco HyperFlex HX Data Platform

## Enterprise-class storage features

The data platform has all the features that you would expect of an enterprise shared-storage system, eliminating the need to configure and maintain complex Fibre Channel storage networks and devices. The platform simplifies operations and helps ensure data availability. Enterprise-class storage features include the following:

- **Replication** stripes and replicates data across the cluster so that data availability is not affected if single or multiple components fail (depending on the replication factor configured).
- **Native replication** transfers cluster data to local or remote clusters for backup or disaster-recovery purposes. Data transfer is based on a many-to-many relationship between the local and remote cluster nodes: each node participates in the data transfer, minimizing the performance impact.

## Cisco HyperFlex HX-Series Nodes



Cisco HyperFlex HX220c M5  
All-NVMe, hybrid, Edge, and  
All Flash Nodes



Cisco HyperFlex HX240c M5,  
hybrid, All Flash, LFF, and  
Edge Nodes



Cisco HyperFlex HX240c M5  
Short Depth Edge Node

- **Synchronous replication**, or stretch clusters, places copies of data in two locations at the same time, enabling cluster failover with a zero RPO and very short recovery time objective (RTO).
- **Data deduplication** is always on, helping reduce storage requirements in virtualization clusters in which multiple operating system instances in client virtual machines result in large amounts of replicated data.
- **Data compression** further reduces storage requirements, lowering costs. The log-structured file system is designed to store variable-sized blocks, reducing internal fragmentation.
- **Encryption** protects your data at rest with self-encrypting drives combined with enterprise key management software.
- **Thin provisioning** allows large data volumes to be created without requiring storage to support them until the need arises, simplifying data volume growth and making storage a “pay as you grow” proposition.
- **Fast, space-efficient clones** rapidly replicate storage volumes so that virtual machines can be replicated simply through metadata operations, with actual data copied only for write operations.
- **Data protection API** enables enterprise backup tools to access data volumes for consistent, per-virtual-machine backup operations.

## Engineered on Cisco UCS technology

Cisco UCS, the foundation for Cisco HyperFlex systems, is built with a single point of management and connectivity for the entire system. The system is designed as a single virtual blade server chassis that can span multiple chassis and racks of blade and rack server-based nodes. Cisco thus is in the unique position of being able to deliver a hyperconverged solution that can incorporate blade and rack systems in its architecture, offering greater flexibility than any other solution. Because Cisco develops the servers upon which Cisco HyperFlex nodes are based, you can count on having the latest processor technology available to speed your performance.

You can optimize your system with the amount of computing and storage capacity that you need by changing the ratio of CPU-intensive Cisco UCS blade and rack servers to storage-intensive Cisco HyperFlex capacity nodes. You can further optimize environments requiring graphics acceleration such as virtual desktop infrastructure by configuring GPUs into your nodes, speeding graphics rendering and helping to deliver a superior user experience.

### Cisco HyperFlex HX-Series Nodes

A cluster requires a minimum of three homogeneous nodes (with disk storage). 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. Each node is equipped with at least one high-performance SSD drive for data caching and rapid acknowledgment of write requests.

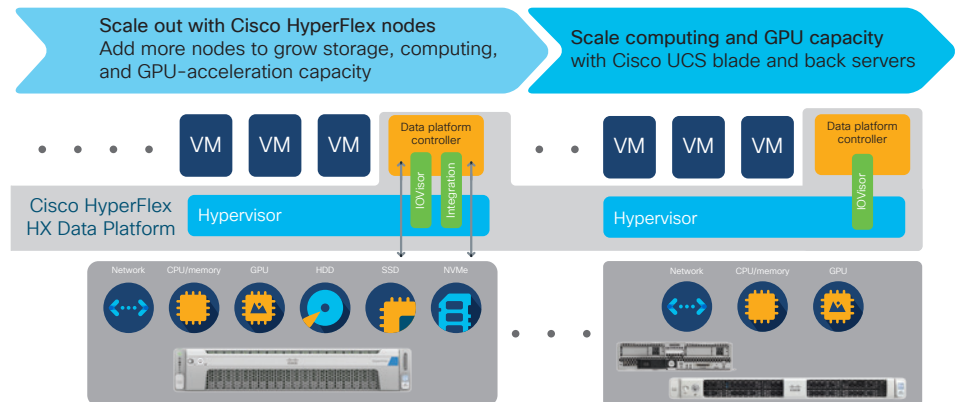


Figure 7 Scaling options for Cisco HyperFlex clusters

- **Hybrid nodes** use SSD drives for caching and HDDs for the capacity layer.
- **All-flash nodes** use SSD drives or NVMe storage for caching, and SSD drives for the capacity layer.
- **All-NVMe nodes** use NVMe storage for both caching and the capacity layer.

Hybrid, all-flash, and NVMe options are available in three types of nodes:

- **Cisco HyperFlex HX220c M5 Hybrid, All Flash, and All NVMe Nodes** balance up to 8 drives or NVMe devices in a 2-socket, 1-rack-unit (1RU) package ideal for small-footprint environments.
- **Cisco HyperFlex HX240c M5 Hybrid and All Flash Nodes** balance high disk capacity (up to 23 drives) in a 2-socket, 2RU package ideal for storage-intensive applications.
- **Cisco HyperFlex Edge Nodes** are designed to work in simplified two- to four-node clusters using existing 1 and 10 Gigabit Ethernet networks. These nodes are available in 1RU hybrid and all-flash, 2RU hybrid and all flash, and 2RU short-depth form factors. Cisco HyperFlex Edge systems offer the same easy deployment and management as all Cisco HyperFlex systems.

### Independent resource scaling with Cisco UCS servers

If you need more computing power or GPU acceleration, you can add Cisco UCS servers to your cluster to increase the ratio of computing power to storage. These servers become computing- and GPU-only nodes that participate in the data platform layer but with no local cache or storage. You need to have a minimum of three HX-Series nodes in your cluster, and you can add Cisco UCS servers up to the point where you have two Cisco UCS servers for each HX-Series node in the cluster. Supported servers include the Cisco UCS B200 M4 and M5 Blade Server and the Cisco UCS C220 and C240 M4 and M5 Rack Servers (Figure 7).

### Hyperconvergence at the network edge

We took a comprehensive look at what is needed at the next frontier in the digital transformation, and we developed the capability to deploy virtually

## Solving the split-brain problem

Traditional two-node edge deployments require a witness node in order to prevent a split-brain situation. This situation can arise if a network failure interrupts communication between the two nodes and each one keeps running. This can result in inconsistent data because two nodes think they own the distributed file system. To prevent this problem, a witness node joins one of the two nodes in the event of a failure. This creates a quorum, a majority vote of two out of three nodes, that allows the cluster to continue to operate without corrupting data.

This adds enormous cost and complexity by having to maintain a virtual or physical node and high-speed connectivity between every edge site and each witness node. Just keeping track of the relationship between edge nodes and witnesses can be daunting when hundreds of remote sites are managed.

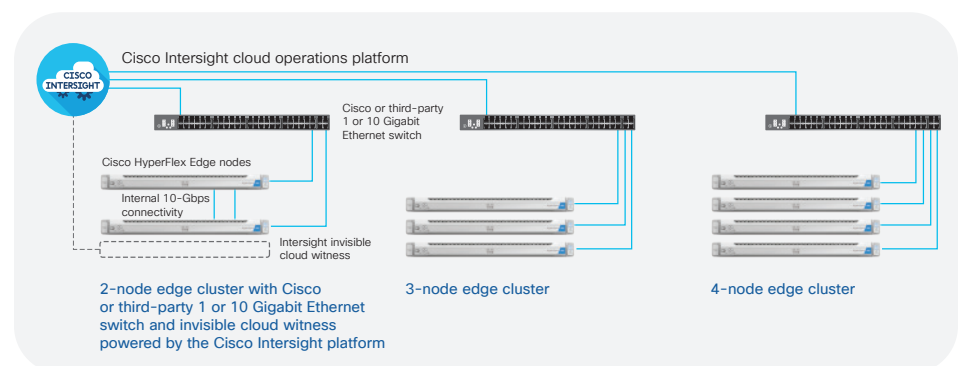
We solve this problem with the **Intersight invisible cloud witness**. It is automatically deployed in the cloud when you deploy a two-node edge cluster. It is simple and doesn't need your attention or require high-speed networking to the edge location. This simple solution enables you to deploy edge clusters with massive scale.

anywhere with massive scale. The combination of Cisco HyperFlex Edge solutions and Cisco Intersight management helps you extend the simplicity and efficiency of your hyperconverged infrastructure from your core data center to the edge with consistent policy-based enforcement and cloud-powered management. Now you can have an enterprise-class edge platform that powers the growing requirements in branch offices and remote sites while enabling new IoT and intelligent services at the network edge.

One factor giving our solutions a competitive edge is that we manage the entire hardware and software stack, so Cisco HyperFlex Edge nodes can be installed by a technician who needs only connect power and network cables. The entire hardware and software stack can then be installed and the cluster deployed remotely through Cisco Intersight. The Cisco Intersight platform gives you end-to-end lifecycle management of edge sites, and can deploy multiple sites identically, in parallel, with full control over the entire stack including firmware, hypervisor, and data platform. This gives you automatic, worldwide reach that facilitates massive scale.

Ongoing management and monitoring can identify failures and automatically capture log files for rapid root-cause analysis. Enhanced security at the edge includes automated deployment of Cisco SD-WAN solutions to securely tie your edge systems into your enterprise data center. Enhancements to native replication support edge locations with capabilities for both local and remote replication to support use cases described in "Native Replication" on [page 12](#).

Cisco HyperFlex Edge configurations can use Cisco or third-party 1 or 10 Gigabit Ethernet switches for connectivity (Figure 8). New in the 4.0 platform release is support for a wider range of edge cluster sizes: two, three, and four nodes. Two-node clusters also can use the nodes' built-in 10 Gigabit Ethernet LAN on motherboard (LOM) to interconnect the two nodes at high speed regardless of the local switch bandwidth.



**Figure 8** Cisco HyperFlex Edge provides scalable and cost-optimized solutions for anywhere deployment.

## Cisco Intersight cloud operations platform

This SaaS interface is available through the cloud or an optional data-center-hosted management appliance. Intersight gives you instant access to all of your clusters—including installation capabilities—regardless of where they are located. This means that you can ship nodes to the facility of your choosing and install them remotely. High-level resource inventory and status is provided by the Intersight dashboard, with drill-down into specific clusters and nodes. New proactive optimization recommendations are derived from telemetry data from each of your clusters.

### Connect to external shared storage

Many organizations use external shared storage in existing environments. Your Cisco HyperFlex systems can connect to that shared storage directly through the fabric interconnects. You can connect through 8- or 16-Gbps Fibre Channel or up to 40-Gbps Small Computer System Interface over IP (iSCSI). With this capability you can:

- **Boot and run virtual machines** stored on the shared storage system
- **Migrate virtual machines** to your more scalable hyperconverged cluster
- **Use shared storage** for backing up your existing environment

Fibre Channel storage can be connected directly to each hypervisor with separate Fibre Channel interfaces that are configured through software on the Cisco UCS virtual interface cards (VICs) in each node. These interfaces allow the cluster to be configured to follow the hypervisor vendor's recommended best practices for traffic separation by creating separate network interfaces for each type of traffic, configured through software.

### Powered by Intel Xeon Scalable processors

Each Cisco HyperFlex M5 node is powered by two Intel Xeon Scalable processors. These processors deliver significantly improved performance and can serve a much wider range of application needs than prior servers. The family delivers highly robust capabilities with outstanding performance, security, and agility. The CPUs provide top-of-the-line memory channel performance and include three Intel UltraPath Interconnect (UPI) links across the sockets for improved scalability and intercore data flow. Each HX-Series node includes a range of processor choices so that you can best serve your application requirements.

### Manage your HyperFlex clusters

Bringing you complete infrastructure automation, Cisco UCS management detects any component plugged into the system, making it self-aware and self-integrating. With the system itself able to adapt to changes in hardware configuration, you need only a few mouse clicks to incorporate new servers into a cluster. Every aspect of a node's identity, configuration, and connectivity is set through software, increasing efficiency and security and reducing deployment time.

Cisco HyperFlex systems integrate easily into existing environments and operational processes and are ready to support DevOps processes and support for your local Kubernetes environment. The management API enables integration into higher-level management tools from Cisco and more than a dozen independent software vendors (ISVs). Tools include Microsoft System Center Virtual Machine Manager, Hyper-V Manager, and the VMware vSphere plug-in.

## For more information

- [Cisco HyperFlex systems](#)
- [Cisco HyperFlex Services](#)

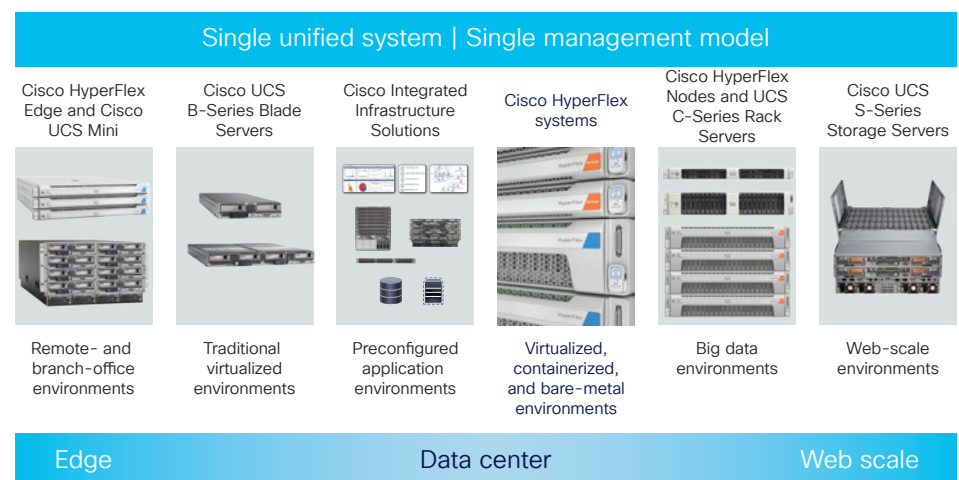
## Conclusion

With Cisco HyperFlex systems, we deliver a complete, next-generation hyperconverged solution that you can use for a wide range of applications and in any location, including your enterprise data center, remote locations, private cloud, and hybrid cloud environments.

We 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 pay-as-you-grow economics of the cloud—but with the benefits of on-premises infrastructure.

### Part of our broad data center strategy

As part of our overall data center vision, Cisco Intersight with Intersight Kubernetes Service helps you deploy, optimize, and manage a hybrid cloud environment with Cisco HyperFlex systems as your on-premises foundation (Figure 9). When you choose Cisco HyperFlex systems, you take your organization beyond a point-product solution, putting your business on a path to a more agile, adaptable, and efficient future.



**Figure 9** A single management model supports the entire Cisco product portfolio.