Cisco HyperFlex Systems

Cisco HyperFlex Anywhere

You need infrastructure that can follow your data and increase the speed of your business regardless of where it takes place. Enterprise applications, big data, and deep learning in your core data center. Virtualized and containerized applications in your private cloud or reaching out into one of many public clouds. Edge locations that put servers closer to your customers and IoT infrastructure in remote offices, branch offices, and retail and industrial sites. Cisco HyperFlex™ systems with Intel® Xeon® Scalable processors deliver hyperconvergence with power and simplicity for any application, on any cloud, anywhere. Engineered on the Cisco Unified Computing System™ (Cisco UCS®), Cisco HyperFlex systems deliver the agility, scalability, and pay-as-you-grow economics of the cloud with the benefits of multisite, distributed computing at global scale.

· **Any application:** Cisco HyperFlex systems support virtualized and containerized deployment with multiple hypervisors and have been tested and validated for numerous enterprise applications.

· **Any cloud:** The platform includes tools for application performance monitoring, application placement, and cloud mobility so that you can design how to deploy your applications and place them wherever your business needs dictate.

· **Anywhere:** You can achieve true global scale with the simplicity of hyperconverged infrastructure that reaches to your network edge. Cloud-based deployment, management, and monitoring scales through templates that make deploying hundreds of remote sites as simple as deploying a single one.

Cisco HyperFlex systems with Intel Xeon Scalable processors

Contents

The challenge: A changing landscape

Computing happens where data resides, and organizations everywhere are recognizing the need for multisite, distributed computing in order to meet the growing requirements in branch office, remote sites, the Internet of Things (IoT), and other intelligent services at the edge. Today you must support an even wider set of application deployment models virtually anywhere the data resides:

- **Traditional enterprise applications** demand data center support for some monolithic components but also for those that are transitioning to a more scalable, cloudlike model for components such as presentation layers.

- **Big data and analytics** need scale-out architectures with large amounts of high-performance storage. Data scientists need massive amounts of GPU acceleration for artificial intelligence (AI) and machine learning (ML) processes. Data enters AI and ML workflows at the edge, works its way to the data center, and eventually drives inferencing engines back in the network edge.

- **Cloud-native architectures** based on microservices demand both virtualized and container-based deployment. Your developers are creating these applications so they can run across multiple clouds with automatic deployment and scaling.

- **Distributed applications** need to be dispersed to where users and data reside. IoT applications use edge locations to acquire and clean data before forwarding the useful parts to the core data center. Remote, branch office, retail, and industrial locations need always-on computing even if the core data center is unavailable. Point-of-sale, video surveillance analysis, virtual desktop, and inventory management are just a few edge applications where IT organizations need to deploy to sometimes hundreds of remote sites.

Reducing complexity is the only way to support both existing and new applications, ensure service delivery, maintain control over data, place computing at remote sites, and attain necessary performance. This need is fueling demand for hyperconverged solutions because they are among the best ways to deliver the wide range of support that IT organizations must provide.

But you can't solve these challenges with just any hyperconverged infrastructure. You need solutions that are:

- **Ready for any application.** You need support for virtualized and containerized software, support for multiple hypervisors, and support for GPU acceleration. Support for enterprise applications is a must.

- **Ready for any cloud.** Today's developers are already creating microservices that deploy literally anywhere. You need tools that help you deploy all of your applications across multicloud environments with analytics to ensure cost-effective choices.

We 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 2020, more than 50% of enterprises will run mission-critical, containerized cloud-native applications in production, up from less than 5% today."

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.

- **Ready to be deployed anywhere.** Organizations with expanding edge applications need to simplify deployment, management, and monitoring so that they can meet aggressive cost goals despite massive distributed scale.

The solution: Cisco HyperFlex systems

To help you meet the challenge of deploying any application, in any cloud, anywhere, we introduce Cisco HyperFlex systems, an adaptive system to power any application anywhere with the simplicity of hyperconvergence. Cisco HyperFlex systems deliver flexibility to support multiple hypervisors, containerized environments, multicloud services, and edge deployment to efficiently and economically deploy, monitor, and manage today's applications.

We designed Cisco HyperFlex systems, powered by Intel Xeon Scalable processors, as a next-generation platform that could adapt to meet new information technology challenges as they evolve. Indeed, the first challenge met by the platform was to reduce cost in the data center core as business organizations rise to meet the challenge of digital transformation (Figure 1). Recognizing the rise of containerized applications and microservices, we
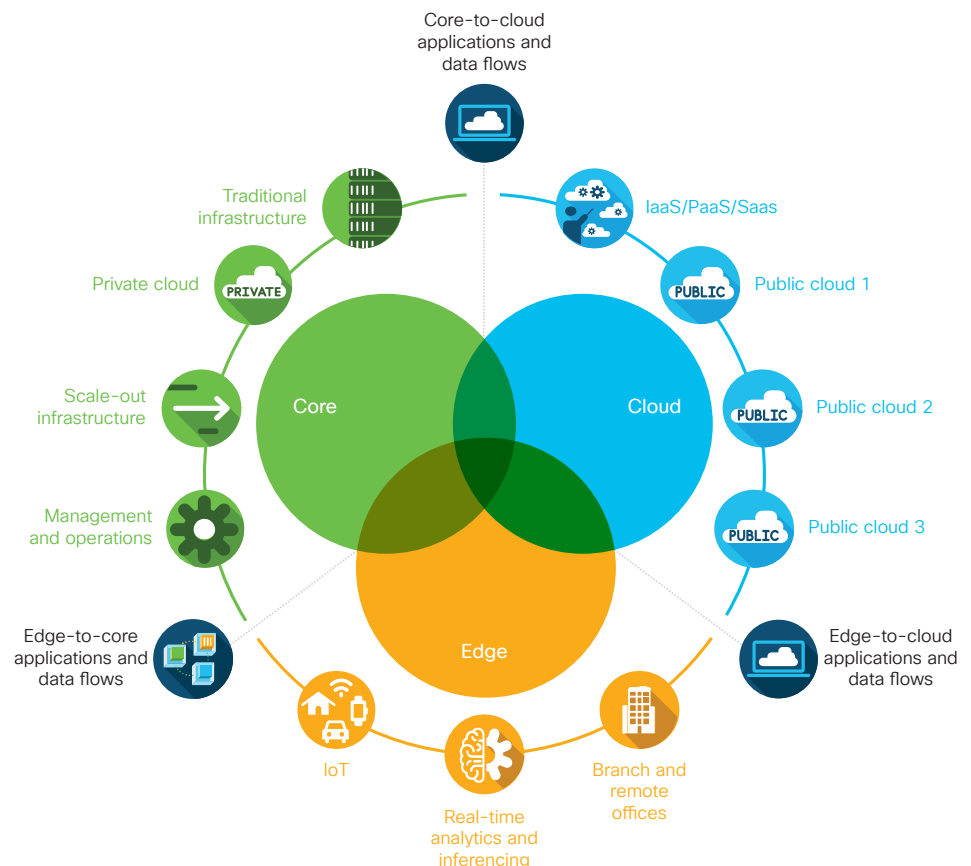

Figure 1. Cisco HyperFlex systems have evolved to support data center core, multicloud, and edge use cases

simplified support for cloud-native applications with the Cisco HyperFlex multicloud platform. The next frontier, the push to place computing wherever data and users reside, has prompted the engineering of Cisco HyperFlex Anywhere to meet the unique challenges of deploying hyperconverged environments for multisite, distributed computing at global scale.

A key difference between Cisco HyperFlex systems and competitive products is that since we produce integrated hardware and software nodes, we are free to innovate across the entire stack. This allows us to deliver capabilities that others cannot, including a unique hardware data compression accelerator, full bare-metal deployment and software installation managed from the cloud, and more GPU acceleration in virtual desktop environments or AI and ML inferencing applications. In data center environments, the platform includes an integrated network fabric and powerful data optimization capabilities that unlock the full potential of hyperconvergence for a wider range of workloads and use cases.

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. You harness these resources with your choice of centralized management tools: Cisco Intersight™ cloud-based management as a service, Cisco HyperFlex Connect, Microsoft System Center Virtual
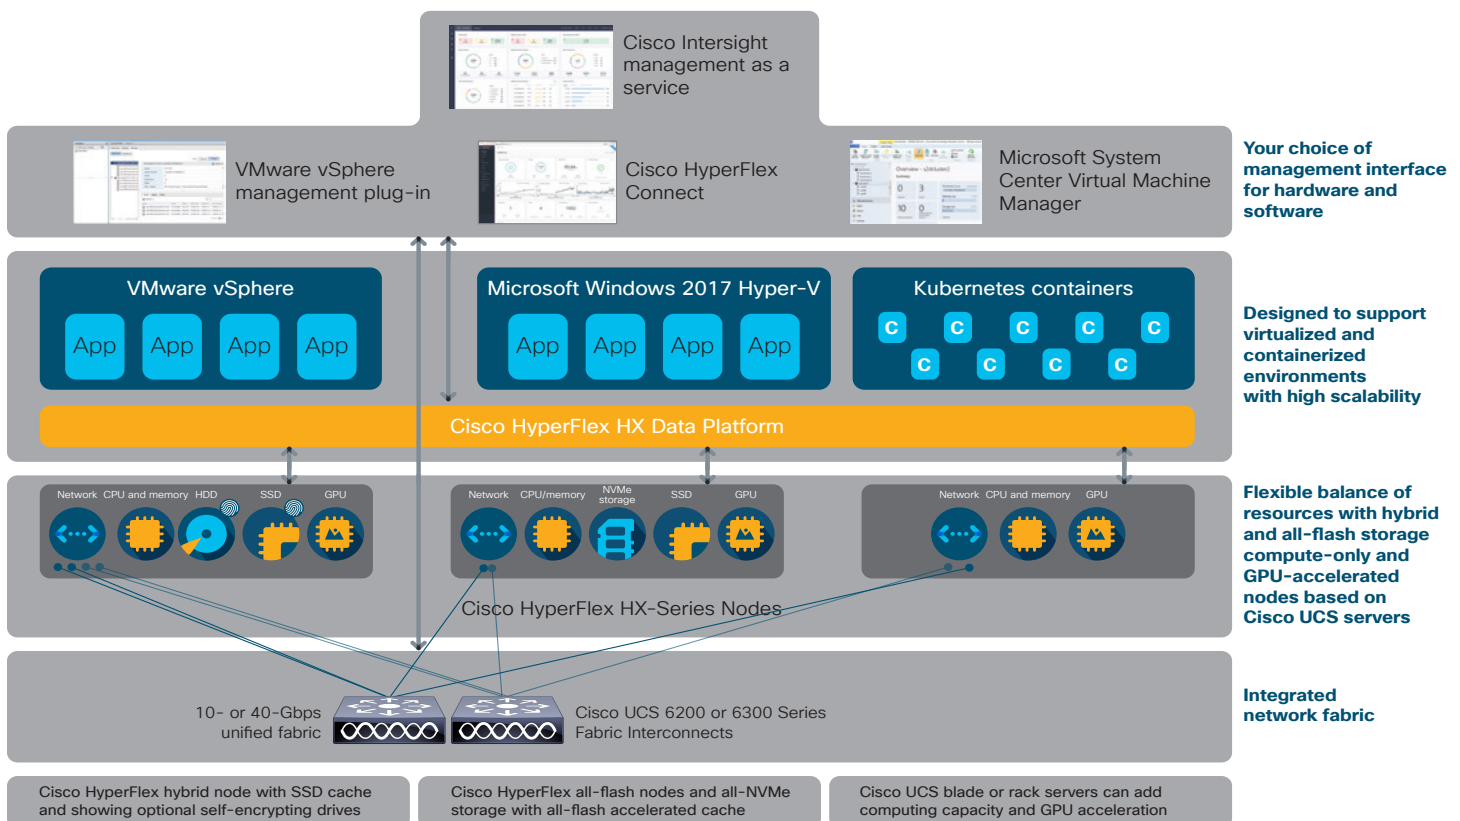
Figure 2. Cisco HyperFlex systems support virtualized and containerized applications and let you choose the exact combination of resources to power your enterprise applications

What's new?

Platform version 4.0 implements Cisco HyperFlex Anywhere capabilities:

- **Flexibility and scale at the edge.** 2-, 3-, and 4-node edge configurations support a wider set of use cases. Template-based lights-out deployment, configuration, management, and monitoring speeds anywhere deployment at scale.

- **Cisco Intersight invisible cloud witness.** Automatic, cloud-based witness for 2-node clusters eliminates the complexity of configuring and maintaining a witness node for each edge site.

- **Enhanced Cisco Intersight management as a service.** End-to-end lifecycle management lets you install, configure, manage, and monitor with a worldwide reach. We integrate the entire hardware stack for lights-out, zero-touch deployment; no other vendor offers this capability. Full-stack provisioning and upgrades keep firmware, hypervisor, and data platform revisions at the level you specify. Parallel and heterogeneous deployment lets you handle massive scale. So do cluster profiles that make it as easy to deploy hundreds of sites as it is to deploy a single one. Connected Cisco Technical Assistance Center (TAC) can automatically initiate support cases based on cloud-based monitoring.

Machine Manager, Hyper-V Manager, or a VMware vSphere plug-in. Cisco HyperFlex systems integrate into the data center you have today without creating an island of IT resources. You can deploy Cisco HyperFlex systems wherever you need them, from central data center environments to remote locations and edge-computing environments (Figure 2).

Complete end-to-end solution

We offer the first hyperconverged platform that is designed as an end-to-end software-defined infrastructure that eliminates the compromises found in first-generation products. 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 multicloud 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** with the powerful Cisco HyperFlex HX Data Platform software
- **Software-defined networking** with Cisco UCS fabric that integrates smoothly with Cisco® Application Centric Infrastructure (Cisco ACI™)
- **Cloud-based management** with Cisco Intersight management as a service, and multicloud container support from Cisco Container Platform for HyperFlex

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 in an hour or less and that scales resources independently to closely match your application resource needs.

Ready for your enterprise applications

The power and performance of Cisco HyperFlex M5 nodes 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. And with flexible configurations, you can deploy Cisco HyperFlex systems in your enterprise data center, to create private clouds, or for edge computing.

Now you can use hyperconvergence for enterprise databases such as Oracle Database and Microsoft SQL Server. The system's performance and high storage capacity makes it suitable for big data applications such as Splunk. High memory capacity and industry-leading performance has resulted in SAP HANA certification. Optional GPU acceleration helps Cisco HyperFlex systems provide a smooth user experience in virtual desktop infrastructure (VDI) environments. Container support allows you to deploy containers in

What's new? (continued)

- **All-NVMe nodes.** These deliver the highest performance for mission-critical data center workloads. We provide an architectural performance edge with NVMe drives connected directly to the CPU rather than through a latency-inducing PCIe switch. Intel Optane™ SSDs also connect to the PCIe bus to accelerate caching for even greater performance than NVMe drives alone.

- **Cisco HyperFlex Acceleration Engine.** Improves performance and efficiency in the data center with faster and lower-latency data compression operations. More compression makes more efficient use of storage resources.

- **Inferencing at the edge.** You can perform deep learning on GPU-only nodes in the data center and drive inferencing at the edge with up to two NVIDIA® Tesla® T4 and P6 GPUs in edge nodes and up to six NVIDIA Tesla GPUs in Cisco HyperFlex HX240c nodes.

your own private cloud and multicloud environments so that you can write once and deploy anywhere. GPU acceleration speeds AI and ML software for model training, machine learning, and inference generation. Our solution is ready to serve your VMware vSphere and Microsoft Windows Server Hyper-V virtualized environments and Internet infrastructure services such as Microsoft Exchange.

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 built to be self-aware and self-integrating so that when the system's fabric interconnects notice a new component attached, that component is automatically incorporated into the cluster. Rather than requiring you to configure each element in the system manually through a variety of element managers, Cisco UCS was built so that every aspect of server personality, configuration, and connectivity can be set through software.

Your choice of management tools

A high level of automation is controlled through the Cisco UCS management API by higher-level tools that allow you to configure a cluster in minutes, regardless of location, with no risk of configuration creep or noncompliant settings. You can choose the approach that best suits your needs.

- **Cisco Intersight management as a service.** This cloud-based interface 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 Cisco Intersight dashboards. A recommendation engine helps you proactively respond to impending issues such as the need to scale capacity. Connected Cisco TAC integration automatically opens a support case when errors are detected and can even initiate shipment of replacement parts such as disk drives. 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. Drill down into data platform operations and you have exactly the same control as Cisco HyperFlex Connect provides—except that you do not have to host management software and you always have the most up-to-date versions. Storage analytic capabilities can track and monitor compute, network, and storage configurations and provide proactive optimization recommendations.

- **Cisco HyperFlex Connect.** This intuitive, data-center-hosted HTML-5-based management tool is device independent, giving you access to all cluster features through any device, anywhere. All of the cluster data platform features (such as logical availability zones and

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 hypervisor and data platform 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.

native replication) can be managed through this interface. Cluster profiles enable you to quickly replicate cluster configurations for fast and accurate deployment. Single-click, full-stack upgrades allow you to keep your firmware, data platform, and hypervisor versions up to date easily.

- **Hypervisor management integration.** Cisco HyperFlex systems management integrates with both Microsoft Hyper-V and VMware vCenter. This level of management integration enables you to manage the lifecycle of virtual machines—including storage management through the data platform—without leaving the interface that your administrators are probably already familiar with. These management tools include Microsoft System Center Virtual Machine Manager (SCVMM), Microsoft Hyper-V Manager, and a VMware vCenter plug-in.

- **Container and multicloud depoyment.** You can build and deploy microservices based on Kubernetes containers with Cisco Container Platform for HyperFlex and Red Hat OpenShift Container Platform. With click-of-the-mouse simplicity, you can write your applications once and deploy anywhere. Cisco Container Platform for HyperFlex allows you to seamlessly deploy Kubernetes clusters in your own private cloud and the Google public cloud

Other tools make use of the management API to support your application lifecycle. Cisco UCS Director can implement end-to-end application lifecycle orchestration and automation. Cisco ONE™ Enterprise Cloud Suite can turn your cluster into an on-premises cloud service with smooth migration of workloads to and from public cloud facilities.

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.

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

More efficient

Our solution is more efficent because of the following:

- **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). If you already have enterprise shared storage, you can connect it directly to your cluster to run virtual machines from it, migrate the storage to your cluster, or use the storage for backup operations.

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 real-time, always-on 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 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 Cisco 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
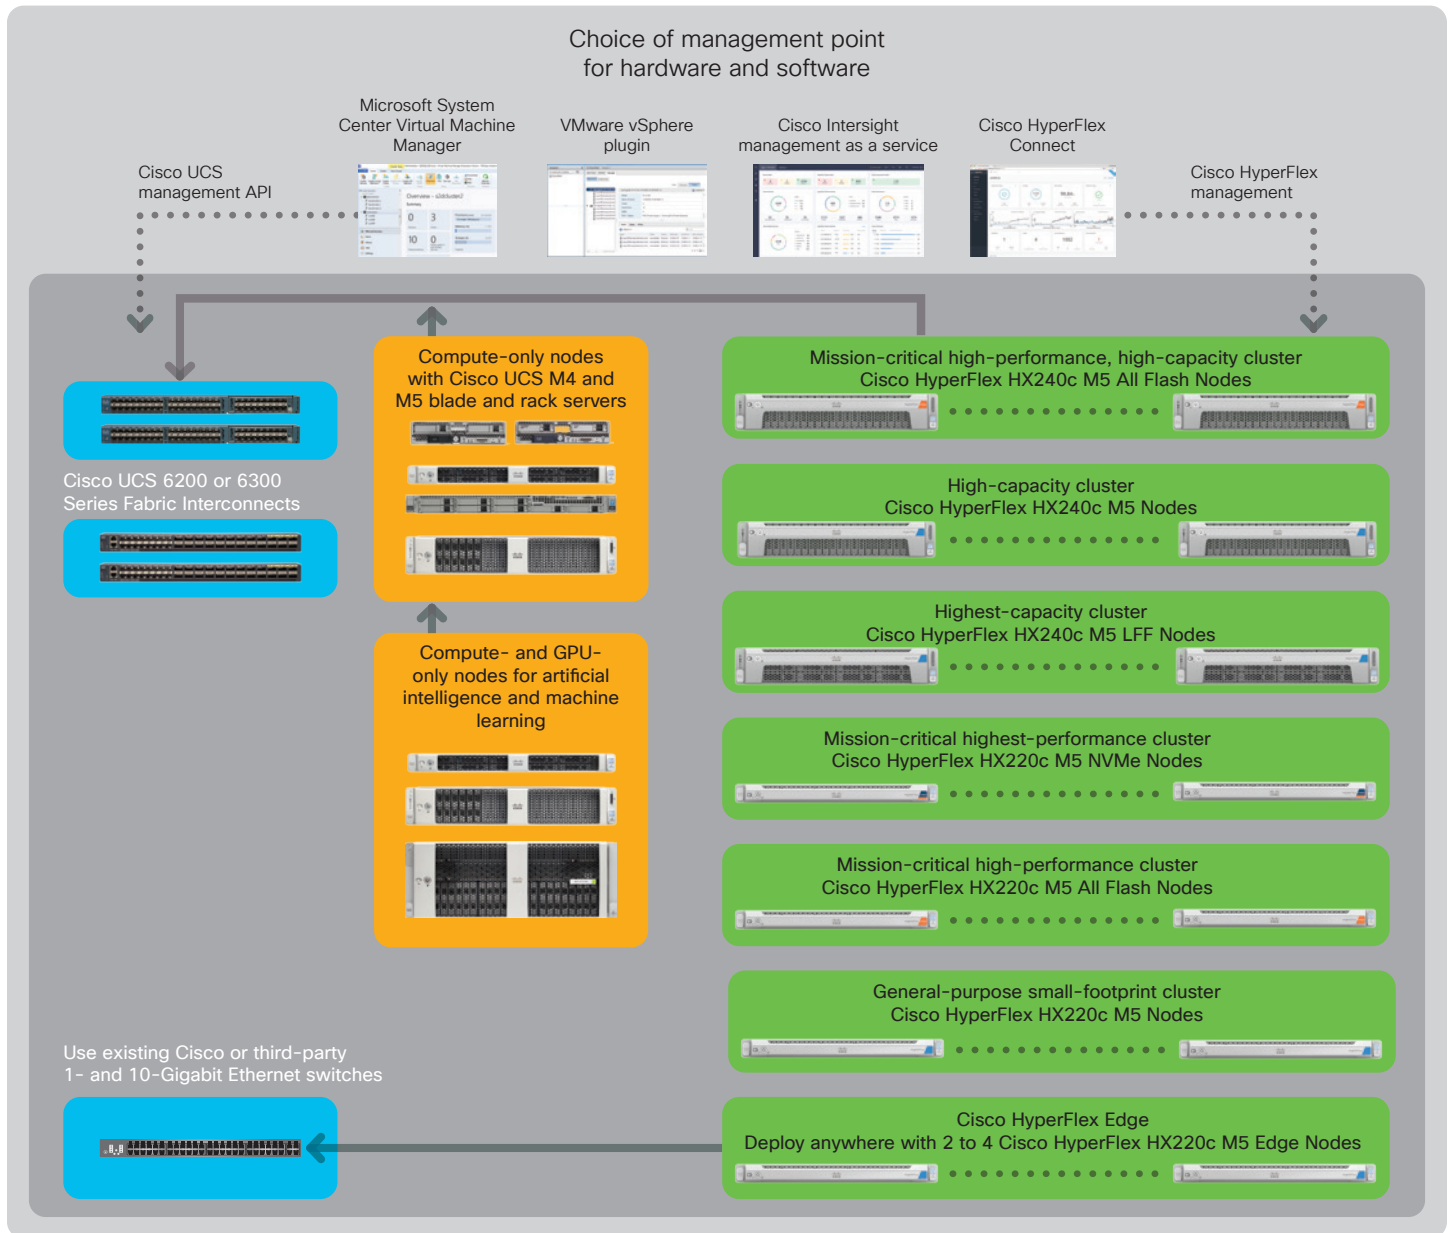
Figure 3. Cisco HyperFlex systems product family

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 6200 or 6300 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)

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.

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.

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 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

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.

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.

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

- **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.

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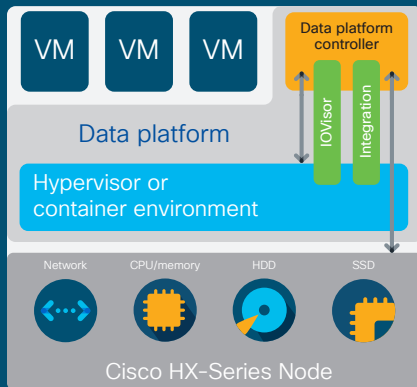
Figure 4. Cisco HyperFlex Data Platform Controller plugs into the hypervisor in each node

the local and remote clusters participating in the replication activity, the overhead imposed on each cluster node is minimal.

· **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.
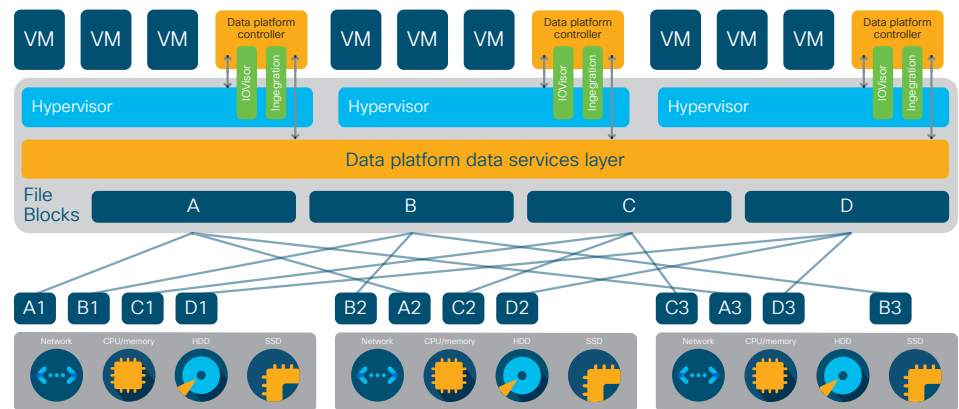
Cisco HyperFlex HX Data Platform controller

A Cisco HyperFlex Data Platform controller resides on each node and implements a distributed 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.
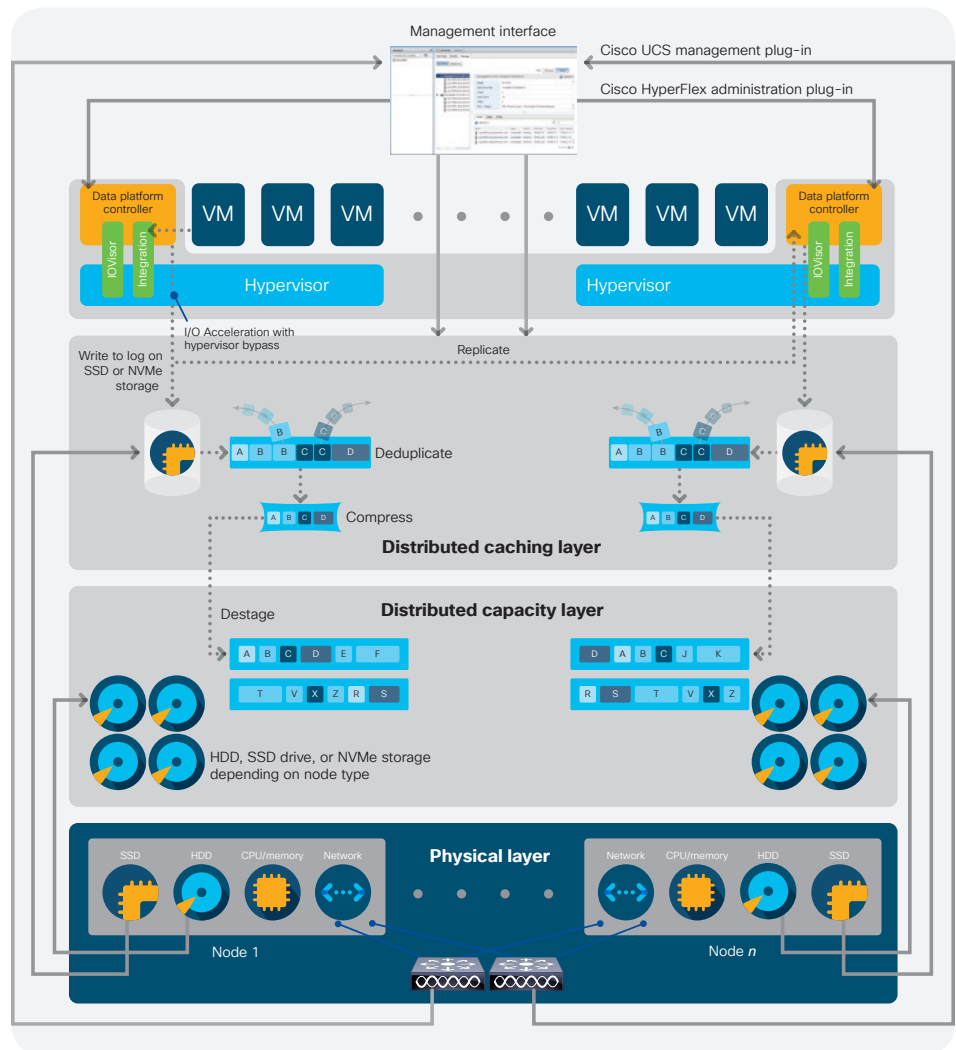
Figure 6. Data write flow through the Cisco HyperFlex HX Data Platform

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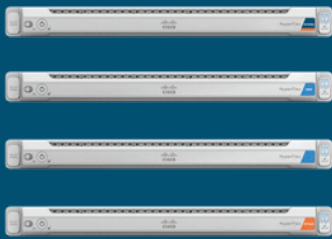
Cisco HyperFlex HX220c M5
All-NVMe, hybrid, Edge, and
All-Flash Nodes

Cisco HyperFlex HX240c M5,
All-Flash, and All-LFF Nodes

- **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.
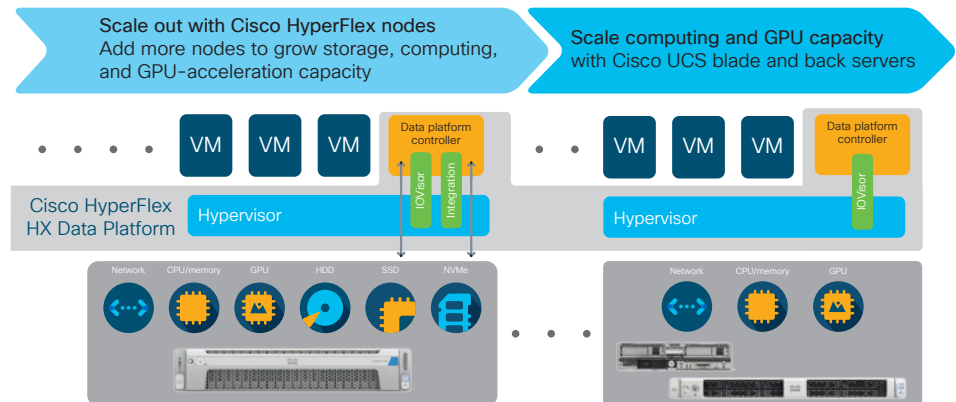
Figure 7. Scaling options for Cisco HyperFlex clusters

- **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 HX220c M5 Edge Nodes** are designed to work in simplified two- to four-node clusters using existing 1 and 10 Gigabit Ethernet networks. Cisco HyperFlex Edge systems offer the same easy deployment and management as all Cisco HyperFlex systems.

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 anywhere with massive scale. The combination of Cisco HyperFlex Edge

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.

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 cloud-based management. 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.

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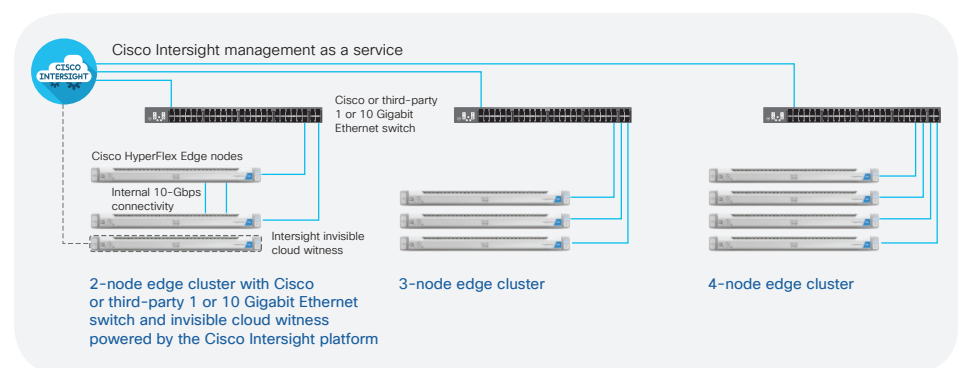
Figure 8. Cisco HyperFlex Edge provides scalable and cost-optimized solutions for anywhere deployment.

Cisco Intersight management as a service

This [cloud-based interface](#) accesses the Cisco UCS management APIs to give you instant access to all of your clusters—including installation capabilities—regardless of where they are located. This means that you can ship your 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 provided by a hosted version of Cisco HyperFlex Connect. 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 arrange 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.

Cisco UCS management

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. With Cisco UCS service profiles, 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. Cisco UCS management's API enables integration into higher-level management tools from Cisco and more than a dozen independent software vendors (ISVs). Tools include cloud-based Cisco Intersight management as a service, Cisco HyperFlex Connect interface, Microsoft SCVMM, Hyper-V Manager, and the VMware vSphere plug-in.

Tools that also use the Cisco UCS Management APIs include monitoring and analysis tools such as VMware vRealize Operations Manager and vCenter, other deployment and configuration tools, and service orchestration tools such as VMware vRealize Orchestrator. When Cisco HyperFlex systems are managed with Cisco UCS Director or Cisco ONE Enterprise Cloud Suite, they can be managed as infrastructure as a service (IaaS) or even as part of a hybrid cloud along with other Cisco and third-party hardware.

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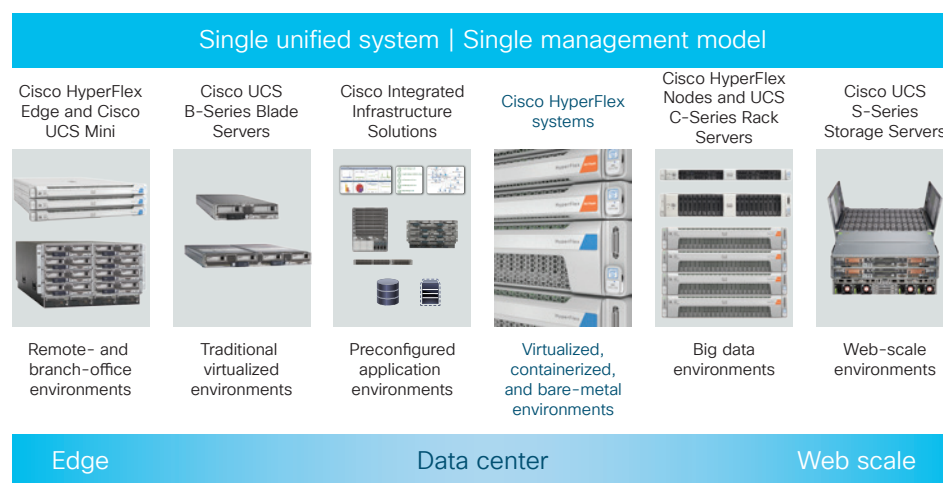
Figure 9. A single management model supports the entire Cisco product portfolio.

Part of our broad data center strategy

As part of our overall data center vision, integration with Cisco Enterprise Cloud Suite will put you on the road to a hybrid cloud environment (Figure 9). Cisco UCS Director enables you to use your Cisco HyperFlex systems to deliver IaaS and create a hybrid cloud. And when you need to augment capacity to handle periodic peaks, Enterprise Cloud Suite will help you use a

For more information

- Cisco HyperFlex systems
- Cisco HyperFlex Services

policy-based approach for engaging public cloud services. Cisco Container Platform helps you develop and deploy in multicloud environments. 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.