



# Deploy Hyperconvergence anywhere your business goes

## In this issue

---

Research From Gartner Market Trends: Sharpen Your HCI Edge Approaches With the Right Value Propositions	2
Cisco HyperFlex HX Data Platform: a new level of storage optimization	12





Research From Gartner

## Market Trends: Sharpen Your HCI Edge Approaches With the Right Value Propositions

Edge computing, a high-growth segment, is driving providers to explore new growth opportunities by positioning HCI. Technology product marketers seeking to grow agile infrastructure opportunities can focus on edge computing use cases to boost visibility and sales.

### Key Findings

- HCI offers the capability to provide multiple benefits, including optimizing IT operations for network edge locations, which is motivating HCI vendors to start to expand exploring edge computing as a lucrative use case.
- Existing edge computing infrastructure is typically a multilocation infrastructure investment for an enterprise, coupled with challenges around centralized monitoring, remote administration, on-site IT support and latency.
- HCI providers primarily focus go-to-market programs on the data center rather than computing at the edge.

### Recommendations

Technology product marketers seeking to grow agile infrastructure opportunities should:

- Position HCI as one of their viable edge-specific infrastructure alternatives by highlighting HCI's relevance in various use cases and the typical three major edge-centric deployment scenarios explained in this document.



- Create an ecosystem of various providers by partnering for monitoring tools, remote IT management providers, system integrators (SIs) and smart racks.
- Highlight all of the benefits of HCI at the edge such as remote management, rapid deployment and centralized monitoring in the messaging.

### Strategic Planning Assumptions

By 2025, as a result of digital business projects, 75% of enterprise-generated data will be created and processed outside the traditional, centralized data center or cloud, an increase from less than 10% in 2018.

By 2021, 40% of large enterprises will be integrating edge computing principles into their projects, up from less than 1% in 2017.

### Introduction

Edge computing creates a location-specific solution to meet a range of use cases. Customer requirements include mitigating latency, data locality, data sovereignty, security, performance, cost and reducing the impact of application latency due to networking. This kind of solution facilitates data processing at or near the source of data generation. Edge computing serves as the decentralized extension of the campus network, cellular network, data center network or the cloud. It provides near-real-time insights and facilitates localized actions. Data processing capabilities such as event filtering via a rule engine, complex-event stream processing or postprocess analytics will be executed near the source of data generation. Edge computing infrastructure can transmit data to the cloud or any data center to provide advanced insights, historical analysis, archiving or higher-level aggregation.

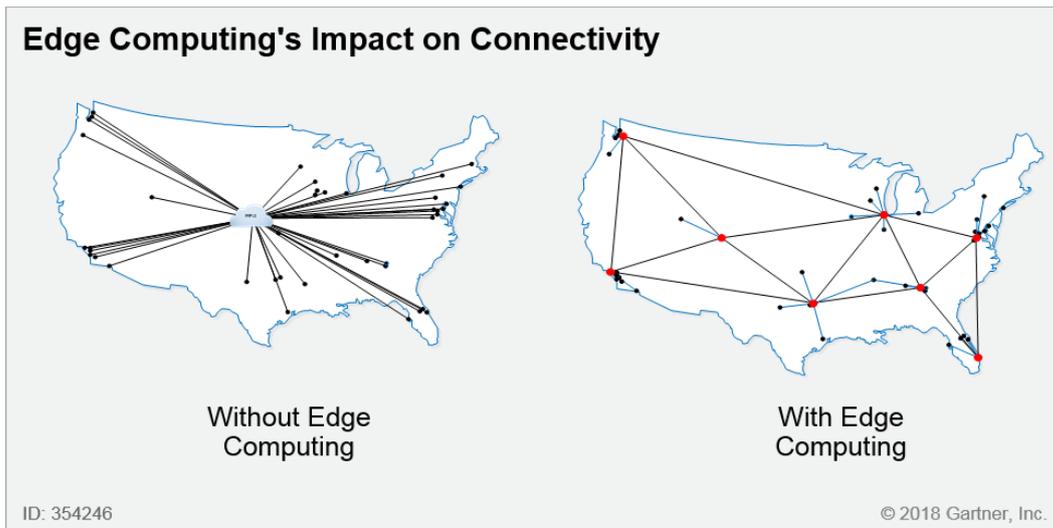
An edge computing solution can be mobile (such as on vehicles), fixed (such as part of a building management solution, manufacturing plant or offshore oil rig) or a mixture of both (such as in hospitals and medical settings). Edge computing solutions can be highly verticalized. For example, edge computing in the context of CSPs would be covered under the concept of “mobile edge computing,” where cellular base stations are treated as edge locations. These edge location host servers push the application data to subscribers in the location. This can be treated as an extension of the CSP’s core network. Various other use cases, such as IoT, augmented reality (AR), virtual reality (VR), drones, autonomous vehicles (AVs) and health services, will be driving the setting up of edge locations hosting edge computing.

Figure 1 demonstrates connectivity that can be simplified by rearchitecting with the addition of computing at edge locations. Connectivity not only reduces the latency and cost, but also delivers higher resiliency.

Hyperconverged infrastructure (HCI) is a software-centric architecture that tightly integrates compute, storage and virtualization resources in a single system that usually consists of industry-standard hardware. It dispenses with the need for a regular storage area network (SAN). HCI is delivered to the customers as either an appliance, or as “software.” HCI is being perceived by end users as an offering with compelling cost advantages over traditional SAN-based or network-attached-storage-based approaches. This has contributed to its rapid growth.

Helping drive IT innovation, HCI supports a growing number of enterprise workloads. Alongside this growth, in the first six months of 2018, Gartner client inquiries about edge computing have been

**Figure 1. Edge Computing's Impact on Connectivity**



Source: Gartner (August 2018)

accelerating rapidly, up more than 20 times, over the full-year 2017. Edge computing offers opportunities as an emerging, but high-growth, use case that can boost the growth of HCI.

Edge computing offers new opportunities for innovation, savings and improved business data acquisition. However, it is a multifaceted technology with much to be affirmed in terms of technologies and services, leaving many businesses with compatibility and service quandaries to consider. Technology product marketers must align the many facets and products to produce a compelling and uncomplicated proposition highlighting the benefits of HCI and edge computing working in harmony.

Figure 2 shows results of polling conducted at Gartner's IT Infrastructure, Operations Management & Data Center Conference 2017 in Las Vegas. It shows 27% of the respondents have already adopted

edge computing, and 54% understand the benefits of the concept and are keen to analyze more and evaluate before moving to deployment mode. Technology product marketers clearly need to invest in building customer awareness efforts around edge computing, with specific focus on HCI.

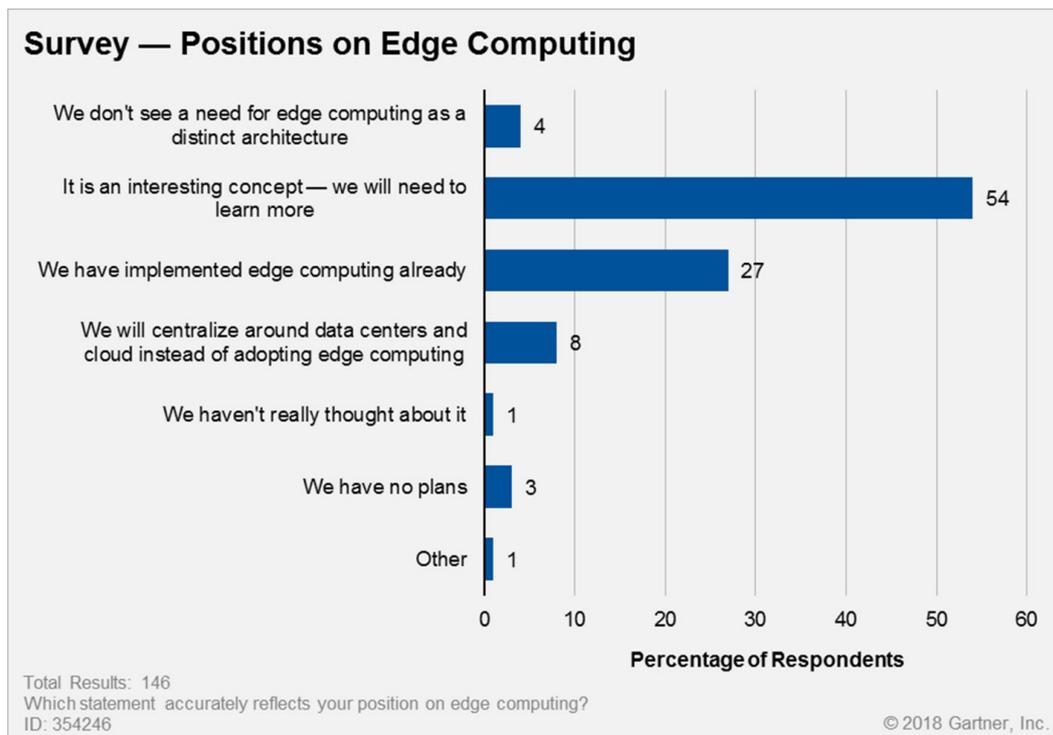
Technology product marketers at HCI providers should proactively position edge computing as an emerging use case for HCI, and build an ecosystem by engaging in new partnerships.

### Market Trend

## The Market for Edge Computing Can Help Drive HCI Adoption

HCI growth opportunities continue as the market transitions to reference architecture and software-driven growth models. A range of use cases is helping drive adoption.

**Figure 2. Survey — Positions on Edge Computing**



Source: Gartner (August 2018)

In partnership edge computing, opportunities will be widespread as businesses understand the emerging use cases and benefits that the business can derive from it. However, deployment requires joined-up thinking by end users and vendors to deliver the right solutions for edge computing.

Businesses such as shipping, transportation, manufacturing, multilocation retail, closed-circuit TV surveillance and others that leverage IoT, need a computing solution at the edge locations. The need to invest in the big infrastructure such as the data center, cooling and power backup sometime delay either the decision or the implementation. In this kind of scenario, technology product marketers should build messaging on the need for HCI at the edge location as a better fit to the requirement.

Similar edge opportunities exist in the remote office/branch office (ROBO) seen in most of the verticals, but primarily in banking, pharmacy, supply chain, retail, energy and utilities, and government, with broader edge computing deployments providing further opportunities. As more applications that need localization functionality for delivering outcomes increases, the need for highly reduced network latency, faster access to data and storage increases, as these applications need to be deployed at the edge. Technology product marketers should use this opportunity and focus messaging on the benefits of HCI, on how HCI would reduce some of the current challenges at the edge.

When it comes to the mentioned verticals, we expect to see HCI replacing traditional three-tier

architecture technologies consisting of storage, server and network appliances, for select use cases. HCI, combined with the right set of tools for edge computing, can help improve remote management services and operations, covering large distributed networks with low-cost, fault-tolerant and powerful systems keeping cost in control and boosting operational efficiency.

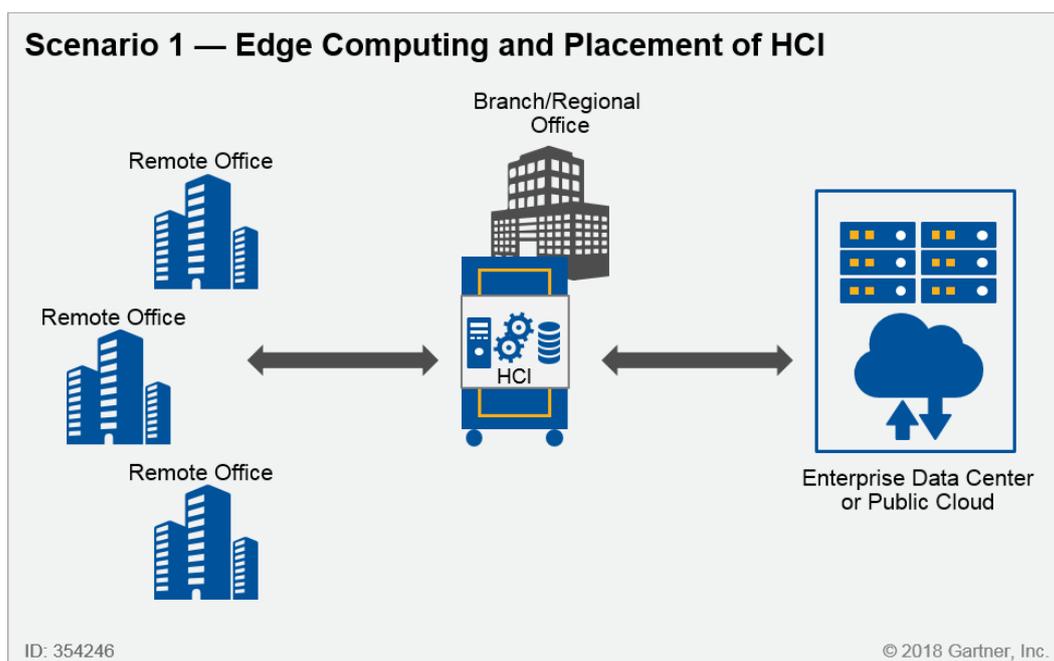
Based on the way the computing is placed and consumed by users, there are typically three deployment scenarios of edge computing. Figures 3 through 5 depict the use cases and provide information on placement of HCI for these use cases.

**Scenario 1** — In this scenario, there are few users at remote offices. The requirement for the business applications for the day-to-day tasks would not

require many resources, typically, a client/server model. The computing is placed at the branch/regional office where the business had made the investment in computing. Because there are fewer resources managing IT in this location, this location becomes an ideal for investment in entry-level HCI nodes. This type of scenario is mostly seen in verticals such as insurance, utilities and government.

**Scenario 2** — This is a scenario typical to an IoT setup. In this scenario, the logs generated from the assets are aggregated at the gateway, and then the logs are, ideally, moved to the edge servers. Many users run data processing tasks on these edge servers. If there is a need from the business, the data is either copied/moved to the central data center and/or public cloud. As the need for computing and storage is needed at the location

**Figure 3. Scenario 1 — Edge Computing and Placement of HCI**



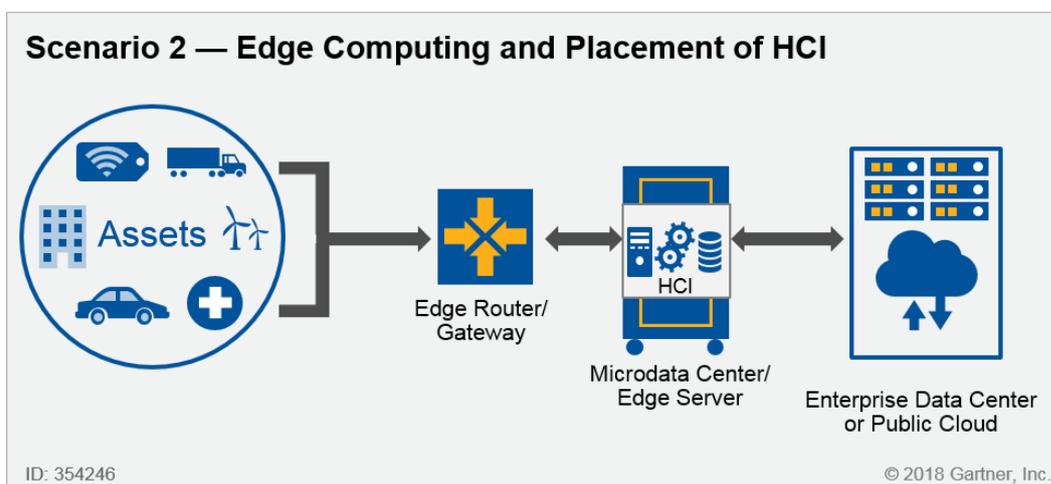
where the edge servers are placed, it becomes ideal for customers to invest in entry-to-midrange HCI solutions and place them in-lieu of the edge servers. This scenario is applicable to all verticals that have adopted IoT. In some situations where the edge gateway and edge servers are placed in one location, connectivity to the public cloud and enterprise data center happens from that location.

**Scenario 3** — In this scenario, the computing is meant to run at each of the branch/remote offices as the remote offices hold more users. The computing is hosted out of this location. As there are fewer resources managing IT in the edge location, the IT at the edge location becomes an ideal for investment in HCI. This scenario is typically seen in verticals such as retail, shipping, banking, insurance, energy, healthcare, government and manufacturing.

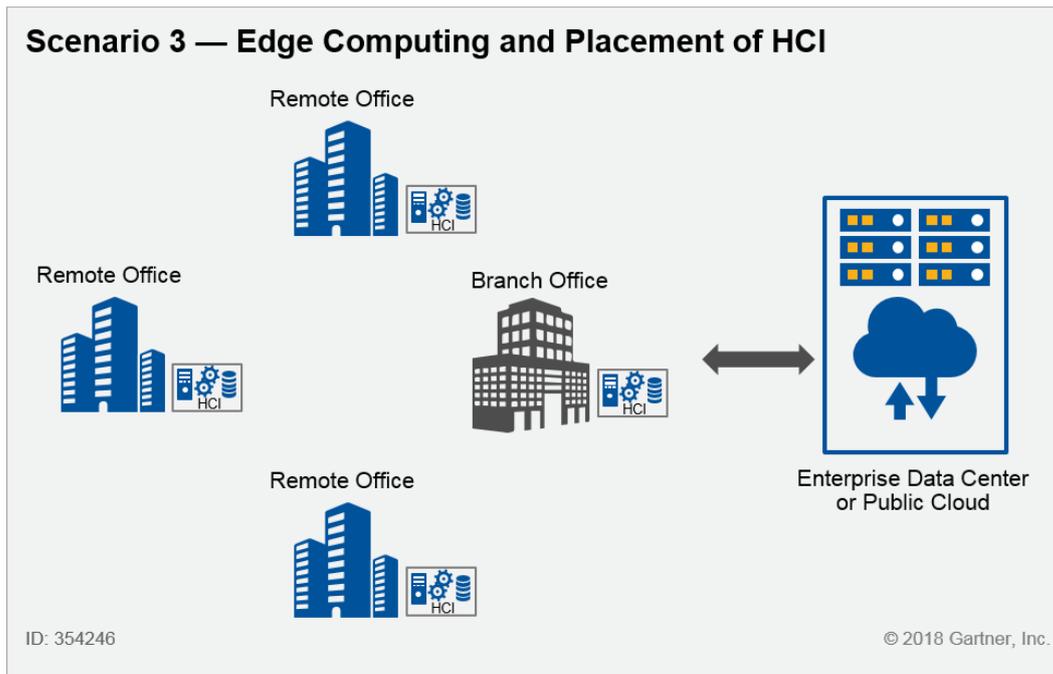
Across these three scenarios, a need for lower total cost of ownership through power efficiency, lack of space in remote areas, flexible configuration requirements, remote management and rapid deployments all help make HCI on the edge an attractive proposition. Technology product marketers with HCI providers should proactively engage with the solution and product teams to identify the right HCI product for each of these scenarios. Some of key use cases that will possibly drive edge computing and make positioning of HCI relevant are AR, VR, drones, AV and content. The following will provide some insights to some of the use cases:

- **Content** — The oldest form of edge computing is content delivery network (CDN). With online video and other content industries growing with the increased adoption of the internet, the need for multiple CDN locations will grow.

**Figure 4. Scenario 2 — Edge Computing and Placement of HCI**



**Figure 5. Scenario 3 — Edge Computing and Placement of HCI**



Source: Gartner (August 2018)

- **Augmented Reality** — Localization for the AR consumer is critical, and this will be especially true when retail and entertainment begin augmenting their advertising and other services with AR. This will impact bandwidth, latency, processing and storage. This market's growth would be affected if edge wasn't already a reality and economical.
- **Autonomous Vehicles** — AVs are predicted to be the future of transportation and delivery, but are in their early stage of development. The investments and number of vehicle manufacturers building/testing prototypes are increasing day by day. In most of the cases, the edge computing needed for an AV is already in

the vehicle. To make the AVs more effective, it makes sense for the vehicle to send data related to traffic conditions, road conditions, local events and visibility around corners to a nearby edge location. The same information would be relayed to other AVs in that location for the other AVs to take necessary action.

However, currently, we see minimal movement from HCI providers to position their solutions for edge computing requirements. Some may need to make technology changes in their products, especially when it comes to entry-level node configuration and solutions. Technology product marketers need to position HCI as one of the viable edge-specific infrastructure alternatives by highlighting HCI's relevance in the multiple use cases and align solutions to the deployment scenarios.

As investment in edge continues to grow, we expect that the related technology trends such as cloud computing, data analytics and IoT will align with HCI. This will provide technology product marketers the opportunity to ride the hype with clear marketing solutions that cut through businesses' confusion. Technology product marketers need to highlight the ecosystem of various providers by partnering for monitoring tools, remote IT management providers, SIs and smart racks. IT infrastructure and network monitoring tool integration will help customers get detailed insight into the health of the IT and network infrastructure at the edge location. This will provide advantage and assistance to the remote IT management providers. Partnering with SIs will exploit the large opportunities that they are bidding on. As the need to leverage an alternative to building a data center at the edge location, smart racks or micro data centers are possible options today. Technology product marketers need to partner with the providers of the smart racks and micro data centers, and build joint go-to-market plans.

## Buyer Trend

### Buyers Seek Low-Cost Ways to Add Compute Power to Remote Locations

While edge computing and HCI will be the preserve of large and midsize enterprises for the immediate future, the potential cost savings and emerging use cases could see smaller buyers adopting the technology. Gartner's "Cool Vendors in Edge Computing, 2018" highlights vendors providing real-time edge analytics engines, IoT security, edge-focused AI solutions for services such as facial recognition,

among others. These applications will help buyers understand the possibility for HCI and edge.

For now, the information points to adoption around key HCI use cases, with technology product marketers having to focus on a vertical-specific approach, tailoring products to each vertical market, with each having different needs. Key to any solution will be intense and intelligent use of data analytics. The benefits, including reduced overheads in data traffic, improved responsiveness and better access to data as an end result, provide operational efficiencies, cost reductions and asset optimization. Technology product marketers need to align the messaging of their HCI to the customer's business goals by highlighting the value of HCI at the edge.

Edge won't be a single market, due to its amorphous nature. With many markets having various edges, the market and growth in solutions will be wide. For example, a growing number of edge computing solutions for the industrial IoT will help producers and factories improve their efficiency, while the edge for consumers will be just as vital, but markedly different in approach.

### Vendors to Watch

The HCI vendors listed in Table 1 have started leveraging benefits of edge computing and ROBO by building use cases around it.

**Table 1. HCI Vendors With Plans for Edge Computing Positioning**

<b>Vendor</b>	<b>Remarks</b>
Cisco	Cisco HyperFlex Edge is specifically designed for edge/ROBO and the system is delivered as a cluster of three hybrid or all-flash nodes, connected over existing Gigabit Ethernet switches. It can manage the central and remote locations by leveraging VMware vSphere plug-in.
DataCore	DataCore Hyper-converged Virtual SAN solution with two servers is typically positioned for ROBO as a use case. However, this provider does not have a targeted strategy for edge computing.
Hewlett Packard Enterprise	HPE has a dedicated Edgeline series of servers that are more-positioned for deployment at remote locations, but are typically fit for locations with extreme conditions. HPE has been carefully positioning SimpliVity for virtualized workloads and typical ROBO scenarios.
Huawei	Huawei leverages FusionCube for its ROBO/edge computing solution. Though the provider has witnessed limited success with ROBO, there is an opportunity for clearer articulation of its overall edge strategy in various messaging targeted for IT leaders.
Nutanix	With the launch of one- and two-node clusters in 2018, Nutanix showcased its intention to cater to the needs for edge/remote office requirements. These one- and two-node clusters can support all of a remote office's local applications and services such as Dynamic Host Configuration Protocol, DNS, file and print. This one- and two-node cluster comes prepackaged with AHV, Nutanix's native hypervisor, and Prism, an infrastructure management tool. In addition, Nutanix's software version can be used on any industry-standard hardware.
Pivot3	Pivot3 Edge Appliance Acuity X3 is a one-rack form factor with all the features of its enterprise HCI. This product uses Nonvolatile Memory Express flash architecture and it includes Pivot3's policy-based Dynamic Quality of Service engine. The provider has seen some success in this space.
VMware	VMware announced a remote office HCI product in early 2018, a smaller variant of its vSAN storage-based HCI. It can run in as few as two servers. It includes vSphere and VMware Pulse IoT Center for management and security of all edge systems/gateways and sensors.

## Acronym Key and Glossary Terms

<b>AR</b>	augmented reality
<b>CDN</b>	content delivery network
<b>AV</b>	autonomous vehicle
<b>CSP</b>	communications service provider
<b>HCI</b>	hyperconverged infrastructure
<b>IoT</b>	Internet of Things
<b>ROBO</b>	remote office/branch office
<b>SAN</b>	storage area network
<b>SI</b>	system integrator
<b>VR</b>	virtual reality

## Evidence

- Vendor briefings
- Vendor website
- Inquiries
- Gartner's IT Infrastructure, Operations Management & Data Center Conference 2017 survey, in Las Vegas



# Cisco HyperFlex HX Data Platform: a new level of storage optimization

## Hyperconvergence anywhere

Cisco HyperFlex™ Edge brings the robust feature set and simplicity of Cisco HyperFlex systems to your edge environments with a flexible, scalable, low-cost, centrally managed solution that can be deployed and maintained with massive scale.



## The power of Cisco HyperFlex systems

We designed Cisco HyperFlex systems as a next-generation platform capable of adapting to meet new challenges as IT organizations face them. The first challenge was to help reduce cost and complexity in the data center core. Next we simplified support for cloud-native applications. Cisco HyperFlex Edge was designed to address the next frontier: placing computing anywhere customers and data reside



## Designed for anywhere computing

Datacenters are growing increasingly decentralized as enterprise applications become dynamic between core datacenters, public cloud services, and now edge computing in branch and remote locations to bring IT resources closer to users, IoT devices, and organizational touchpoints. This move to the edge in particular poses new challenges not faced by IT teams previously. Budgets for remote sites are tight. Edge locations have different requirements and must flexibly scale up or down. They must be highly



resilient and be able to operate independently, and without support of the core data center. They must install and operate with minimal or no onsite IT staff. They must be centrally deployed, managed, and maintained. And they must be able to support new inferencing applications with GPU acceleration as needed.

Cisco HyperFlex Edge helps you meet the unique challenges of deploying simplified, hyperconverged environments for multisite, distributed computing with global scale. It incorporates key features optimized to lower cost and reduce space consumption. You can choose clusters with two, three, or four nodes for ease of meeting a wide range of edge-location computing, GPU acceleration, and storage requirements (see Figure 2 on the next page). Cisco HyperFlex Edge locations can use existing Cisco® or third-party 1- and 10-Gbps networks for cluster communication. Two-node clusters can use built-in 10-Gpbs LAN-on-motherboard (LOM) ports for high-speed cluster connectivity even with Gigabit Ethernet switches upstream. These enhancements make Cisco HyperFlex Edge easy and affordable to deploy as a hyperconverged solution in a multitude of edge locations, or even as the sole cluster supporting a small or medium-sized business.

### Deploy and manage with massive scale

The biggest challenge for organizations extending computing to the network edge is handling deployment and management with massive scale. Imagine composing your hyperconverged infrastructure through a single interface that accesses hundreds of clusters at once, with support for installation, inventory management, and day-to-day centralized control. Cisco Intersight™



management as a service automates cluster deployment and management regardless of cluster size or location.

### Deploy

All you need to get started in a remote location is to connect power and network cables to the edge nodes in the remote location. Technicians can do this without any specialized expertise. Once the edge nodes are connected to the Intersight platform through a secure Transport Layer Security (TLS) connection, you claim them in the interface. Then you associate a cluster profile that specifies the complete configuration for the remote site. Installation is fully automated without any manual intervention required. Cluster settings can be embodied in profiles that can direct rapid, consistent deployment without any variance from the standards you set. Cloning tools can be used in the Intersight user interface. If you wish to automate using your own software, the Intersight API can be accessed from a wide variety of scripting tools.



### Configure

To configure Cisco HyperFlex HX Data Platform parameters, all you need to do is click to launch the HyperFlex Connect interface. This enables you to control every aspect of the data platform operation, including creating and mounting data stores, taking snapshots, configuring remote replication, and cloning virtual machines

Configuring two-node clusters is easy. Just install the cluster and Cisco Intersight automatically provides an invisible cloud witness that eliminates the need to set up a witness node in the core data center or hosted location.



Traditional systems require setting up a witness node for each two-node cluster to prevent split-brain situations. This problem occurs when both halves of a cluster think they can continue to operate independently. Traditional witness approaches impede deploying at scale because each witness needs to be installed, maintained, and connected to the edge cluster with high-speed networks. This adds cost and complexity to edge deployments. In contrast, the Intersight invisible cloud witness simply arbitrates which node gets to continue operating in a split-brain situation automatically, without any time-consuming or expensive setup.

### Monitor

The Intersight platform monitors your edge cluster health and provides an intuitive dashboard so that you can easily see the status of all of your nodes worldwide. Just click on any cluster or node to see more status details. The platform is directly connected to the Cisco Technical Assistance Center (Cisco TAC). If Intersight software detects errors that indicate or a hard failure, the Cisco TAC can automatically diagnose the issue by reviewing logs through an artificial intelligence engine that looks for diagnostic signatures.



### Maintain

Cisco Intersight software is adept at maintaining each cluster in the state you desire. The software can maintain an unlimited number of clusters with different hardware and software configurations. Unique in the industry, single-click software updates can update an entire cluster's firmware, hypervisor, and data platform software to the revision levels that you specify. Other vendors can't do this because they



don't integrate the entire hardware and software stack. Cisco's ownership of both these layers helps ensure the most seamless, efficient, and time-optimized upgrades possible on any hyperconverged infrastructure stack.

When a single-click update is initiated, by cross-launching into the HyperFlex Connect interface, you can upgrade multiple clusters in parallel, across multiple sites, without any workload disruption or human intervention needed.

### Stay up to date

Cisco Intersight itself is a containerized, cloud-native application that is continually updated, helping assure that you always have access to the latest, most secure, and most stable version.



### Management options

In addition to Cisco Intersight management, you can also manage your hyperconverged clusters through locally hosted Cisco HyperFlex Connect software or a range of third-party tools and plugins including the VMware vSphere plug-in.



### Next steps

Cisco HyperFlex Edge brings the next generation of hyperconverged solutions beyond your data center and out to the network edge. With access to the same full-featured data platform and management model as you use in your data center systems, extending your reach has never been easier.

### Cisco HyperFlex Edge



- Reduce cost and minimize complexity
- Support always-on applications close to customers and data
- Deploy quickly with massive scale
- Provide resilient computing and storage
- Deploy anywhere, manage centrally



### Cisco IT Testimonial for Hyperflex Edge

#### We're customer zero for HyperFlex

Cisco IT is "customer zero" for new products like HyperFlex. That means that when you receive a new release, we've already put it through its paces for our business and given our feedback to the business unit. The current release of HyperFlex includes our suggestions, such as scaling to 64 nodes per cluster (up from 8 in the original release), an all-flash storage and NVMe cache option, support for containers on VMware and bare metal, and availability zones.

#### Business value of HyperFlex Edge to date

Today we're using HyperFlex Edge for core services in remote offices: Boxborough, Sydney, Bangalore, Dubai, Bedfont



Lakes, and Amsterdam. So far we're seeing the following benefits:

- 4x faster deployment. A 3-node HyperFlex Edge system with VMware ESX takes just four hours to deploy compared to eight hours for three rack servers with third-party storage. "We just cable up HyperFlex, load the installer, and then we're off and running," says John Thomas, senior IT engineer. Saving six hours in each of 75 offices will work out to 450 hours, or 56 days. That's time our staff can redirect to strategic projects. 
- OpEx savings. "We'll eventually spend about 50% less time on management because there is no requirement of multiple compute and storage admins managing the deployment: one person can do it all," says Srikanth Makineni, senior IT engineer. And with built-in storage instead of separate storage, HyperFlex is lowering power and cooling costs. 
- CapEx savings. Compared to three rack servers with third-party storage, CapEx for a 3-node HyperFlex Edge is about 20-28% less. 

We recently launched proofs of concepts for two other use cases: supply chain and unified communications. In supply chain locations, HyperFlex will run quality control and testing software. "Since our manufacturing locations are spread across the globe, we want infrastructure that's cost-effective, quick and easy to deploy, and can be managed from the cloud," says Ashish Nanjiani, senior IT manager. "HyperFlex is just what we were looking for."

## Next steps: centralized management through Intersight



To make managing HyperFlex even simpler we're planning to use Cisco Intersight™, a cloud service that works with all recent Cisco servers: blade, rack, or hyperconverged. With Intersight we'll be able to centrally manage all of our HyperFlex systems instead of logging into each system one by one. "We're looking forward to the convenience of using a single interface to upgrade HyperFlex software, UCS firmware, and—coming soon—VMware ESX," says Bharath Malapaka, IT architect. Managing HyperFlex from the same interface we use to manage UCS blade and rack servers will also make it simpler to track global inventory—for example, to find all servers anywhere in the world that need a firmware upgrade. The more HyperFlex systems we deploy, the more time we'll save.

The unique data demands imposed by applications, particularly those hosted in virtual machines, have resulted in many storage silos. A foundation of Cisco HyperFlex Systems, the HX Data Platform is a purpose-built, high performance, log-structured, scale-out file system that is designed for hyperconverged environments. The data platform's innovations redefine scale-out and distributed storage technology, going beyond the boundaries of first-generation hyperconverged infrastructure, and offer a wide range of enterprise-class data management services.

The HX Data Platform includes these features:

- **Multi hypervisor support** including



Microsoft Windows 2016 Server Hyper-V and VMware vSphere

- **Containerized application support** for persistent container data in Docker containers managed by Kubernetes.
- **Enterprise-class data management features** provide complete lifecycle management and enhanced data protection in distributed storage environments. These features include replication, always-on inline deduplication, always-on inline compression, thin provisioning, instantaneous space-efficient clones, and snapshots.
- Support for hybrid, all-flash, and all-NVMe nodes allows you to choose the right platform configuration based on your capacity, application, performance, and budget requirements.
- Simplified data management integrates storage functions into existing management tools, allowing instant provisioning, cloning, and metadata based snapshots of applications for dramatically simplified daily operations.
- **Improved control** with advanced automation and orchestration capabilities and robust reporting and analytics features delivers improved visibility and insight into IT operations.
- **Improved scalability** with logical availability zones that, when enabled, automatically partition the cluster so that it is more resilient to multiple-node failures.
- **Independent scaling** of Cisco HyperFlex HX-Series nodes, compute-only nodes, and compute- and GPU-only nodes gives you the flexibility to scale your environment based on both your

computing and storage requirements. As you add HX-Series nodes, data is automatically rebalanced across the cluster, without disruption, to take advantage of the new resources.

- **Continuous data optimization** with inline data deduplication and compression increases resource utilization and offers more headroom for data scaling.
- **Dynamic data placement** optimizes performance and resilience by enabling all cluster resources to participate in I/O responsiveness. Hybrid nodes use a combination of solid-state disks (SSDs) for caching and hard-disk drives (HDDs) for the capacity layer. All-flash nodes use SSD or Non-Volatile Memory Express (NVMe) storage for the caching layer and SSDs for the capacity layer. All-NVMe nodes use high-performance NVMe storage for both tiers. This approach helps eliminate storage hotspots and makes the performance capabilities of the cluster available to every virtual machine. If a drive fails, restoration can proceed quickly because the aggregate bandwidth of the remaining components in the cluster can be used to access data.
- **Stretch clusters** synchronously replicate data between two identical clusters to provide continuous operation even if an entire location becomes unavailable.
- **Enterprise data protection** with a highly-available, self-healing architecture supports nondisruptive, rolling upgrades and offers options for call-home and onsite support 24 hours a day, every day.

- **API-based data platform architecture** provides data virtualization flexibility to support existing and new cloud-native data types, including persistent storage for containers.
- **Tools that simplify deployment and operations.** The Cisco HyperFlex Sizer helps you to profile existing environments and estimate sizing for new deployments. Cloud-based installation through Cisco Intersight management as a service allows you to ship nodes to any location and install a cluster from anywhere you have a web browser. Single-click full-stack updates enable you to maintain consistency by updating the node firmware, HX Data Platform software, and hypervisor with a rolling update across the cluster. Cisco Intersight storage analytics will help enable real-time monitoring and predictive intelligence from compute, networking, and storage.

### Hyperconvergence close to you and data



You need infrastructure that can follow your users and data and help scale your business regardless of where computing takes place. Internet of Things (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 edge applications where IT organizations need to deploy to sometimes hundreds of sites.

---

Source: Cisco



## Contact us

To know more about Cisco HyperFlex, Visit:

<https://www.cisco.com/go/hyperflex>

For additional information contact us at:

Call: 000 8000 040 2262 – 9 am – 6 pm IST

Email: [in\\_assistant@cisco.com](mailto:in_assistant@cisco.com)



Cisco (NASDAQ: CSCO) is the worldwide technology leader that has been making the Internet work since 1984. Our people, products, and partners help society securely connect and seize tomorrow's digital opportunity today. Chuck Robbins is the Chief Executive Officer of the \$48.7 billion company. Cisco commenced India operations in 1995. There are seven Sales Offices in the region - New Delhi, Mumbai, Bangalore, Chennai, Pune, Kolkata and Hyderabad. The Cisco Global Development Center is in Bangalore and is the largest outside the US. It houses Cisco's Research and Development (R&D), IT, Services and customer support teams and develops disruptive business models for Cisco to create new go-to-market channels, markets, processes and technologies for emerging markets. The Advanced Global Briefing Center located here showcases Cisco's latest technology solutions and proof-of-concepts.

Deploy Hyperconvergence anywhere your business goes is published by Cisco. Editorial content supplied by Cisco is independent of Gartner analysis. All Gartner research is used with Gartner's permission, and was originally published as part of Gartner's syndicated research service available to all entitled Gartner clients. © 2019 Gartner, Inc. and/or its affiliates. All rights reserved. The use of Gartner research in this publication does not indicate Gartner's endorsement of Cisco's products and/or strategies. Reproduction or distribution of this publication in any form without Gartner's prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. The opinions expressed herein are subject to change without notice. Although Gartner research may include a discussion of related legal issues, Gartner does not provide legal advice or services and its research should not be construed or used as such. Gartner is a public company, and its shareholders may include firms and funds that have financial interests in entities covered in Gartner research. Gartner's Board of Directors may include senior managers of these firms or funds. Gartner research is produced independently by its research organization without input or influence from these firms, funds or their managers. For further information on the independence and integrity of Gartner research, see "Guiding Principles on Independence and Objectivity" on its website.