White Paper

How HCI Is Powering Hybrid Cloud Optimization to Accelerate Digital Transformation

Sponsored by: Cisco and Intel

Stephen Elliot Lucas Mearian Eric Sheppard
July 2021

INTRODUCTION/EXECUTIVE SUMMARY

Digital transformation (DX) represents an embodiment of the benefits we've seen in productivity and efficiency that are now broadly attached to information and communications technology spending. One of the keystones of the modern datacenter and DX is hyperconverged infrastructure (HCI) aligned with advanced management, orchestration, and automation capabilities, which offer operational and life-cycle management benefits and agility similar to a cloud platform for on-premises or edge environments. Storage is fully integrated and virtualized, so it's no longer a repackaging of legacy infrastructure like it once was. Because the nodes are standardized servers, the cluster architecture is easily scaled out, and all the node and cluster resources are pooled together and act as a single, resilient system. In addition, modern applications running across distributed environments demand advanced management that integrates HCI capabilities with development and deployment processes to reduce costs, improve team collaboration, and deliver operational efficiencies.

Along with server and storage consolidation in datacenters, HCI is now used for implementing not only private cloud architectures but also hybrid cloud architectures that combine on-premises private clouds with public cloud services. Many enterprises are rolling out a hybrid cloud strategy from a single-tenant infrastructure and then investigating how to leverage the public cloud from there; the objective is being able to have lower total cost of ownership (TCO), improve interoperability with consistent operations management, and create greater workload mobility. Enterprise IT organizations are demanding management capabilities that enable standardized orchestration and control across a hybrid cloud environment.

HCI is often the basis for deploying containerized environment via a hybrid cloud architecture that allows DevOps, site reliability engineering (SRE), and IT operations teams to optimize control points for cloud-native apps in the public cloud and bring them back on premises for production runs or for cost optimization as the workload scales.

A hybrid cloud model is key to being able to take advantage of one of containerization technology's key attributes: portability. Unlike virtual machines (VMs), containers share a host operating system environment (OSE), whether it's within a virtual machine or a bare metal server running on premises or in the cloud. Containers can migrate among multiple clusters and between private or public clouds regardless of the host OSE. This key container attribute enables cost optimization through migration automation.

A hybrid cloud model also requires management and automation for the life-cycle management of network, compute, and storage devices, Technology executives must have the ability to manage

cloud-native, container-based environments. Deep levels of insight and visibility into Kubernetes clusters, pods, and hosts as well as the ability to bring this information together with other metrics and data are essential to having an end-to-end view of service availability and identifying and resolving problems in cloud-native applications. Increasingly, many of these applications are distributed applications, as they traverse multiple cloud architectures and have integrations with classic or legacy systems and applications.

In support of cloud-native applications, HCI first enables the on-premises deployment of containerized workloads, allowing DevOps teams to more easily spin up Kubernetes clusters and create a continuous integration/continuous delivery (CI/CD) pipeline. Then, HCI in support of hybrid cloud opens DevOps teams to cloud-based, SaaS data services in support of their containerized environments.

SITUATION OVERVIEW

For decades, innovation was centered around driving capital expenditure (capex) improvements.

Capex improvement models were helped along by Moore's law, the number of cores expanding per CPU, hypervisors abstracting the physical hardware, and then basic storage optimization. It was all based on capital expenditures. In the process, the solutions themselves got more complex and needed specially trained personnel to manage, leading to the development of separate centers of expertise and competence such as compute, storage, and networking. These are very often completely different teams.

During that time, for every dollar spent on capex for servers, $0.50 was spent on opex — things like IT staff required to manage the infrastructure and managing facilities, power, and cooling. Over time, primarily led by the cost of hiring specially trained talent to manage the complexity, opex costs continued to increase, eventually overtaking capex expenses. Today, you're spending more on managing and maintaining the infrastructure and the power, cooling, and facilities costs in general than the cost of the infrastructure itself. The expenditure ratio eventually increased to the point where for every dollar spent on capex for server expansion, $4 were spent on maintaining infrastructure — we were spending too much time and money just keeping the lights on.

IT shops are pulled in many directions, and therefore, users are looking for more efficient approaches to their job so that more time can be spent on innovation. In fact, a recent IDC survey revealed IT staff only spent about 20% of their time on innovation. So first we were chasing capex for too long and then we chased opex.

The complexity of the traditional architecture and the organizational structure it requires was no longer able to deliver the speed and agility of decisions and implementations that businesses needed. For example, infrastructure and operations (I&O) teams' relevance and value to the business have increased as revenue-generating services take a central focus in executive decision making. I&O teams are asked to provide tools, processes, analytics, data access, and dashboards to increase operating efficiencies through cross-team collaboration that enables scale, increased automation, workload placement, and faster decision making with the business. Providing highly available services that deliver a great customer experience becomes even more important as executives add more scrutiny to their purchase decisions and demand choice and integrations to avoid the high cost of vendor lock-in. Hybrid cloud management offers I&O teams an opportunity to manage services across

multiple clouds while providing visibility, insight, and actions that can optimize workload placement based on business policies in cost-optimized scenarios.

To put it another way, IT organizations had been spending too much time on datacenter projects and tasks that didn't drive real value to the business, and that became an unacceptable obstacle for companies that are working through digital and business transformation.

Then along came HCI, enabling a dramatic reduction in the total cost of ownership; flexible resource allocation; increased automation, including for infrastructure life-cycle management; and an ability to rapidly and seamlessly scale the infrastructure.

At a very basic level, HCI combines compute, storage, and network elements into a single unit and then scales out to create a distributed system. HCI eliminates Fibre Channel components, storage arrays, and the three-tiered SAN architecture. It also eliminates the need for storage, server, and networking specialists.

HCI creates an abstraction layer between physical assets — server, storage, and networking — and the user. A user can increase or decrease the capacity as if they're simply taking it from or returning it to a preexisting pool.

The building block of HCI — an HCI node — is simply a commercial off-the-shelf server. This is the essence of software-defined infrastructure that abstracts the underlying hardware. All the nodes are integrated by the HCI software and thus act as a single cluster. The software-defined storage (SDS) element also enables automated functions, such as data reduction, backups, replication, and data migration. HCI software can also enable automated, full-stack, life-cycle management.

HCI Use Cases Are Growing

HCI began as a platform for niche workloads, such as VDI and video surveillance capture, especially in remote office/branch office (ROBO) locations. Overtime, however, HCI grew to replace legacy storage and compute environments for several reasons. HCI is less expensive to maintain and grow. It is based on industry-standard servers, making it easier to utilize the latest generation of processors, persistent memory, and flash storage media. HCI allows companies to scale CPU and storage resources simultaneously or independently, depending on needs. This reduces complexity and enables central control using familiar hypervisor APIs. Now one team of IT generalists can manage that infrastructure, and that reduces IT employee training and helps scale existing resources to manage more infrastructure.

An IDC survey released late last year showed organizations are moving away from expensive, proprietary rack and stack systems to embrace an HCI-based horizonal play. The HCI market has grown by double digits each year for the past decade as software-defined storage solutions proved easy to deploy and dramatically reduced TCO.

Automation offered on top of HCI systems and standardized systems is able to reduce overprovisioning and eliminate storage silos, which has positive impacts beyond capex. In fact, these benefits can directly lead to lower costs of power, cooling, and floor space within the datacenter. HCI solutions often integrate management software that automates many of the complex tasks needed during initial deployment while reducing the number of steps required to provision new workloads. The results are improved IT staff productivity and increased agility within the datacenter.

One recent IDC study found participants credited the Cisco HyperFlex HCI platform for enabling them to build out a cost-effective IT infrastructure for their businesses. They noted cost efficiencies compared with both refreshing their legacy environments and using a public cloud solution for equivalent workloads. On average, they are building out and running Cisco HyperFlex HCI at a 31% lower total cost over five years than building an alternative on-premises environment in terms of hardware, warranty, power, and datacenter costs. Study participants not only spent less in deploying the HCI platform than they would with their previous or alternative solution, including public cloud offerings for several interviewed customers, but also saved in terms of licensing and operational costs such as warranty, power, and facilities costs.

HCI products are available as a fully integrated appliance, a certified node, or a software-only offering. The products also offer bundled services such as data deduplication, data protection that includes snapshots and backups, disaster recovery, and WAN optimization. Fully integrated appliances are and will continue to be the most preferred form of procurement of HCI, according to *Worldwide Hyperconverged Infrastructure Forecast, 2020–2024: Edge Represents the Next Wave of Growth* (IDC #US46621420, June 2020). This is due to simplicity and speed of deployment to accelerate ROI and come fully vetted and supported by a single technology vendor.

At the same time, HCI is growing as a flexible, efficient infrastructure. The COVID-19 pandemic forced digitization to the forefront, and companies had to address new business disruptions, adapt to work-at-home business models, and more fully embrace online commerce.

Enterprises and governments around the world have shifted to a major strategic focus on digital transformation, which has helped drive adoption and implementation of new technologies and related services.

Bolstered by the global pandemic, DX investment is expected to grow at a compound annual growth rate (CAGR) of 15.5% from 2020 to 2023, and it is expected to approach $6.8 trillion total investment as companies build on existing strategies.

Digital transformation is increasing the importance of modernizing an infrastructure to handle the new demands. IT teams need to be more relevant and responsive. They can't put a high value on rack and stack. If they do, they're going to be in the way of digital transformation. They need to be able to deliver resources quickly, including VDI, artificial intelligence (AI), and machine learning (ML) on horizontally scalable platforms.

HCI and Hybrid Cloud Drivers and Attributes

As digital transformation efforts were spurred on by the global pandemic, more organizations embraced a hybrid and multicloud architecture to both address their secondary storage needs and deploy new cloud-native applications.

Most organizations anticipate higher HCI investments and many plan on utilizing HCI for their dedicated private cloud, for edge, and to bridge the gap between on premises and public clouds.

While public and private clouds are similar in the basic approach to infrastructure in that they enable flexible provisioning and agility, there are differences in their abilities to address specific requirements; these include the extent to which they allow scalability, access to new technologies, and compliance/control over the resources.

Public cloud offers the greatest potential cost savings through reduction in capex and automated access to the latest technologies, but the public cloud also has its detractions, including concerns about security, compliance, and governance. There are concerns by users around meeting service-level agreements (SLAs), pricing that can grow quickly along with performance requirements, and increased operational costs particularly when it comes to bandwidth consumption. Not all public cloud instances provide equal capabilities, so it is essential to evaluate the performance, capacity, and security capabilities of the underlying infrastructure to ensure it will meet workload requirements. Conversely, private clouds based on HCI offer lower latency, higher performance, and firewall-based security of sensitive assets.

Public cloud best addresses needs when there is a lot of uncertainty around demand, which leads to greater resource variability. Once there is more predictable demand, a private cloud will typically offer a lower TCO than the public cloud.

Along with capex reduction and access to latest technologies, where those instances based on the latest technologies are chosen, public cloud adoption can reduce opex in the forms of infrastructure administration and facilities costs, time to provisioning, and achieving exactly the right scale on demand.

The public cloud, however, can become more expensive than the private cloud when an organization's software infrastructure has achieved a certain scale. High performance or availability SLAs, bandwidth costs (especially when data needs to move between regions), and elasticity (such as in the form of net additional spend on top of committed spend with the public cloud provider) exacerbate the cost of the public cloud.

In those instances, HCI can provide a solid foundation for the hybrid cloud, if it possesses these attributes:

- Consolidation of compute and storage into a single HCI layer running on x86 rack or blade servers, with automated management and flexibility to customize the hardware infrastructure based on workloads
- Consolidation of storage and IP networks into a single virtualized network layer with automated management
- A multipurpose workload engine that provides choice of hypervisors and that is capable of running virtualized, containerized workloads
- A high-performance and robust software-defined storage layer capable of seamless scale out and providing rich data services and protocols
- A SaaS-style control and operations plane that automates infrastructure and platforms management and that provides observability, insights, and optimizations of the workloads, from core to edge to cloud

In addition, organizations need solutions that enable the seamless expansion and migration from private to public and vice versa. This is where hybrid cloud capabilities make sense. It allows IT to have a single operating model and lets the cost, performance, and other strategic factors dictate the location of the workloads.

Many enterprises are rolling out a hybrid cloud strategy from a single-tenant infrastructure and then leveraging that to connect to the public cloud; the objective is being able to have lower TCO, achieve consistent operations management, and create greater workload mobility. In addition, HCI is often the

basis for deploying a containerized environment via a hybrid cloud architecture that allows DevOps to build and test cloud-native apps in the public cloud and bring them back on premises for production workloads.

In fact, hybrid and multicloud is the gold standard for IT architecture and will drive insights, innovation, and growth moving forward. Of enterprises that have adopted cloud technology, 97% use public and private clouds and 92% are doing business with two or more cloud service providers, according to IDC's 3Q20 *Cloud Pulse Survey.*

Infrastructure software investments that manage and orchestrate will be driven partly by the need to bring on-premises datacenter resources and cloud resources together. This hybrid approach is a critical I&O requirement to have the visibility to correlate and optimize across all of the underlying infrastructure, including the application layer. This complexity drives the need for service transparency and the use of advanced analytic models to identify and resolve problems fast to reduce (or prevent) customer impact and the ability to choose the best-fit cloud architecture for each workload to deliver cost optimization and business risk reduction. It also drives the need for tool integrations from multiple vendors, enabling customers to avoid vendor lock-in while providing long-term licensing transparency and flexibility.

In addition, hybrid and multicloud complexity is forcing a rising number of cloud managers to consider new strategies to improve interoperability, portability, and consistency. In 2020, 86% of those managers indicated they will migrate or repatriate some workloads to on premises from public clouds for ease of management, according to IDC's *Cloud Pulse Survey.* A consistent underlying hardware infrastructure improves workload portability for optimal placement across private and public clouds, as well as edge locations.

Another use of hybrid cloud is to manage "bursting," or addressing temporary peaks in IT demand. Workloads can be moved to address processing needs by automatically migrating them from an HCI-based private cloud to a public cloud. The primary advantage of "cloud bursting" is users only pay for the additional resources when there is a demand for them. It also places processing burdens on the public cloud, freeing up local resources for business-critical applications.

Where HCI, the Edge, and the Cloud Meet

Edge computing is the next frontier for HCI systems. It is one of the top areas of strategic IT investments, especially in the light of companywide digital transformation initiatives.

So, what is edge? Let's start with the core, or central datacenter, that can be onsite, at a colocation facility, or can be a cloud service. In contrast, the "edge" is the vast space or intermediary between the data collecting endpoints and the core that supports all key business decision support systems. Depending on the usage, the edge can have a different meaning for each business. An edge device is typically connected to multiple endpoints to aggregate or process data. Those devices can be routers, switches, or various integrated access devices, and they provide an entry point into enterprise or cloud provider's core networks.

While HCI has been used in many ROBO or edge locations for years, edge locations have taken on increased importance due to the rising volume of data from endpoints, combined with the need for additional processing and analytics of this data at these locations.

An "intelligent" edge is a crucial link between the core and endpoints that provides a distributed compute, data persistence, and network aggregation layer and that serves as the intermediary for analytics of collected data.

Edge computing enables delivery and analysis of data and resources to people and systems in a timely fashion. That's because HCI at the edge can process data closest to where it's created — at the ingestion point. HCI at the edge removes the need for costly bandwidth associated with transmitting data from multiple endpoints back to a core datacenter for processing. And processing data at the edge also eliminates the risk of data loss associated with network failures.

A centralized infrastructure approach faces numerous challenges as organizations look to deploy new and emerging use cases for IoT, AI, and ML. Distributing computing resources to edge locations can enable real-time decision making using onsite data analytics on HCI platforms, minimize network bandwidth costs associated with transmitting data to core datacenters, and decrease service downtime by reducing network connections.

Digital transformation is driving a convergence of information technology (IT) and operational technology (OT), and the intersection of those efforts is most often at the edge. Industrial organizations are automating and optimizing production machinery through technology; those systems, however, have traditionally been managed separately from corporate computing resources. As operation teams began experimenting with machine learning and AI in the cloud to assist with automation and optimization, they discovered latency thwarted their real-time applications. Through edge computing, the problem can be solved by training models in the cloud and running them on local infrastructure. Edge computing also enables companies to move away from proprietary systems to an open architecture that can run mixed workloads.

IDC predicts that by 2023, more than 50% of new enterprise IT infrastructure will be deployed at the edge, up from less than 10% in 2020. According to IDC's Worldwide Edge Spending Guide, the edge market will reach $250.6 billion in 2024, with a CAGR of 12.5% over the 2019–2024 period.

In a recent IDC survey, more than two out of three respondents said that their planned 2021 budgets on edge IT spend will either remain the same or increase due to the global pandemic. The current uncertain business environment has challenged organizations to rethink business models and adapt quickly to changing business demand. And cloud-native approaches to edge computing can increase business agility, giving organizations the ability to rapidly modify operations and launch new products and services.

Edge locations have different operating requirements than traditional datacenters. This is reflected in how infrastructure is designed and managed. Edge (ROBO) locations are typically smaller facilities serving a limited number of users. This makes HCI an attractive option, with host configurations typically consisting of two to four hosts.

HCI and edge also meet to enable connectivity to external cloud services. As growth in IoT and mobile applications continues, HCI at the edge becomes home to massive amounts of data outside of the traditional datacenter. The value in consuming, processing, and deriving insight from data at the edge and in real time, in turn, drives demand for the use of artificial intelligence and machine learning there. While public cloud services can be used for that purpose, there can be processing limitations and there are issues around reliability, data protection, latency, and cost. This is pushing companies to deploy edge computing through HCI as a means to process and monetize their data at the edge and remote office and branch office locations.

Operations teams also realized that while latency associated with cloud services could be prohibitive for real-time AI/ML applications, training models could be deployed in the cloud and then run on local infrastructure. Having a unified Kubernetes management plane is key to bridging the gap between training models in the public cloud and deploying those models for inferencing at the edge or on-premises infrastructure. Deploying an AI inference at the edge or taking a model created through deep learning and deploying it on a device at the edge enables the model to process incoming data and look for and identify whatever it has been trained to recognize.

This model also allows companies to move away from closed, proprietary systems to an open architecture that can run mixed workloads. Besides visibility and insight across multiple clouds, infrastructure, and applications, the need for applying analytics to these unified data pools has never been more important.

Analytic capabilities provide more precise triggers to initiate automated workflows and deeper insights into performance patterns. The ability to increase the productivity of existing I&O teams, and make existing resources more efficient, continues to rise as budgets are more tightly managed and environments become more complex. Development teams need to increase the speed and deployment frequency of applications due to agile and DevOps practices. IDC research indicated that by 2022, 50% of DevOps teams will invest in tools to focus on business KPIs (cost, revenue, etc.), and operations will play a larger role in end-to-end app performance and business impact (see *IDC FutureScape: Worldwide Developer and DevOps 2020 Predictions,* IDC #US44636519, October 2019).

Hybrid cloud environments are becoming the de facto architecture in datacenters, offering enterprises an effective way to scale their usage while enabling placement choices for each workload based on its needs and constraints.

Applications that require four or more 9s availability are likely best suited for a private cloud; scale-out applications that may need additional resources quickly are best suited for the public cloud, where there is an infinite amount of infrastructure resources without any up-front infrastructure investment; and analytics workloads may be best suited for a hybrid cloud as AI/ML algorithms comb through large pools of data. Those pools may come from public websites, data warehouses, and private clouds. Database workloads vary greatly in size and compute requirements, so depending on specific needs they may be best suited for a private cloud or a public cloud.

Edge and HCI can also bring other benefits including:

- Addressing the lack of trained IT personnel to manage the infrastructure at ROBO locations
- Supplying ruggedized form factors to address often hostile environments in remote office locations, which can often pose extreme temperatures with no available cooling, vibrations, or even industrial chemicals
- Requiring other means to secure equipment when infrastructure at the edge cannot rely on physical security of the gear behind datacenter walls

As hybrid cloud buildouts take hold, monitoring and managing the systems, data, and services have created more complexity. Users, therefore, are seeking a single management plane through which they can see and control their hybrid environment.

ESSENTIAL GUIDANCE

Realizing the full potential value of a hybrid cloud strategy starts with modernizing and standardizing on software-defined infrastructure. While on-premises infrastructure has traditionally been the most challenging piece of the hybrid cloud architectures due to the demand it places on IT teams to deploy, manage, and maintain, it is also the area that can return the most overall value. HCI has proven to yield business value in supporting traditional datacenter workloads and is evolving into the de facto hybrid cloud foundation in many of today's cloud strategies.

When considering an HCI platform, look for one that can offer a consistent experience regardless of where the workload is running. It should not matter where the user is connected and interacting with the HCI, whether through an API or UI or what you have, the experience should be the same. It also should not matter whether a workload is running on premises, in the cloud, or at the edge.

A consistent underlying hardware infrastructure makes it easier to manage and migrate workloads to a different location as needed. If the on-premises infrastructure looks like the public cloud infrastructure (i.e., they are both scaled out, software defined), then the whole process becomes easier than trying to mix and match an on-premises, dual-controller array with a scale-out, software-defined public cloud layer.

There's a level of decoupling that should happen between the actual physical location of a workload and the experience a user gets when it has deployed and in use. In other words, look for a user experience where the physical infrastructure and the cloud service appear one and the same; resources and services should simply be scaled or reduced using a simple management interface. For example, if you are running a database on premises, there are certain characteristics you will get from that experience. If, for whatever reason, you then want to run that database in the public cloud, it should happen transparently, and the operational and management experience should not change.

Organizations should ensure R&D investments are targeting centralized and holistic management of multiple architectures. The automation of traditional life-cycle management tasks is also becoming a must-have capability for HCI and hybrid cloud environments.

HCI systems should have robust support for container environments. Also expect customers to have a strong affinity between HCI and next-generation flexible consumption models.

There are also so many opportunities for automation and long-term opex savings through an HCI-enabled hybrid cloud infrastructure, and that includes the ability to rethink how IT teams are structured. Restructuring a team can reduce risk. For example, IT operations often have "heroes" inside the datacenter, or those individuals who have great institutional knowledge and know everything about that storage array or network. The expert eventually leaves, and the organization is left helpless, and that is risk.

Integrate performance monitoring and management of traditional/heritage, cloud-native, and container-based applications to deliver transparency and cost optimization across multiple clouds and to reduce customer-impacting downtime.

In addition, the highly automated nature of HCI solutions also helps reduce the risk of downtime associated with common life-cycle management tasks (e.g., firmware upgrades, system refreshes). The scale-out, software-defined nature of HCI solutions helps eliminate the need for complex and risky forklift upgrades, which have become all too common within the datacenter. Many companies leverage

hyperconverged solutions as a way to improve their disaster recovery/high-availability (DR/HA) processes and costs in ways not possible just a few short years ago. HCI solutions also allow users to reduce the number of technology suppliers involved within a full solution, which helps better coordinate patches and upgrades while reducing the number of support calls needed for the solution.

SOLUTION

Keeping in mind that HCI is a subset of the larger converged infrastructure market, IDC notes that the genesis of Cisco's commitment to HCI can be tied to the very early days of infrastructure convergence. Indeed, Cisco is among a small number of technology suppliers responsible for the emergence of datacenter infrastructure convergence, driving an untold amount of savings and operational benefits over the past decade. Through deep partnerships with other suppliers and considerable R&D investments, Cisco has always been, and continues to be, at the forefront of the rapidly growing converged infrastructure market. Organizations around the world have collectively invested more than $100 billion in converged infrastructure from 2012 to 2020.

More than a quarter of this considerably large amount of market value comes from systems that are built with Cisco Unified Computing System (UCS) servers. Cisco has put its considerable experience to good use, leveraging a decade of leadership and learning within infrastructure convergence when designing its HyperFlex portfolio of hyperconverged solutions. HyperFlex is able to take advantage of the regular cadence of gen-over-gen improvements in UCS servers based on Intel platform generations.

Today, Cisco's HyperFlex HCI solutions enable very high levels of efficiency, agility, and resiliency within the datacenter by enabling tight integration of core infrastructure (compute, storage, networking, and system management), increased levels of automation, and simplified life-cycle management. HyperFlex solutions are fully engineered appliances built on its UCS servers (x86) that provide an abstracted pool of capacity, memory, and CPU cores that are used as the foundation for server-centric workloads (e.g., the hypervisor, VMs, and applications) as well as storage-centric workloads (e.g., data persistence, data access, and data management).

Cisco's HyperFlex solutions enable enterprises to deploy applications from core to edge and in multicloud environments. With all-flash and all-NVMe configurations using Intel Optane SSDs for caching, HyperFlex solutions power mission-critical applications and databases in a datacenter and enable multicloud development and deployment of cloud-native apps. Intel Optane persistent memory support enables extremely fast, nonvolatile memory. HyperFlex solutions extend the simplicity of hyperconvergence to the edge at a distributed scale with two-node solutions that utilize a local or cloud-based witness and simplified remote deployment across multiple edge sites from a centralized management pane.

An IDC Business Value study that included interviews with Cisco HyperFlex customers found HyperFlex solutions have been able to better support business operations while delivering more cost-effective, efficient IT operations. Cisco HyperFlex allowed study participants to improve IT agility and performance, reduce risks, and enable IT operations — translating to an average five-year ROI of 452%. In addition, the use of HyperFlex over a five-year period resulted in a 50% reduction in operational cost and a 93% reduction in the time it takes staff to deploy new servers, according to an IDC Business Value study (see *Business Value of Improved Performance and Agility with Cisco HyperFlex,* IDC #US44881119, March 2019).

HyperFlex fully integrates into Cisco's cloud operations platform, Intersight, that provides a full suite of services, including visualization, orchestration, and application and infrastructure optimization. This ensures efficient and consistent management of both on-premises and edge infrastructure (HyperFlex, Cisco UCS servers, and Cisco networking) and across public clouds (e.g., AWS, Microsoft Azure, and GCP) through a unified, SaaS-based management suite.

With Intersight Cloud Orchestrator (ICO), customers also get a single pane of glass management for their virtual machine or container ecosystem. In addition, HyperFlex also provides an out-of-the-box element manager called HX Connect and integrates into a vCenter environment via a plug-in to enable full active management of the cluster in those environments.

In addition, Cisco offers Intersight Workload Optimizer, a SaaS-based solution that helps IT organizations optimize resource usage, control costs, and deliver required performance across a complex landscape of infrastructure, applications, and workloads, including on-premises and cloud deployments. Cisco Intersight Workload Optimizer operates in real time to ensure workloads get resources when needed to optimize overall system operational performance and health. The solution enables continuous placement, resizing, and capacity decisions to achieve optimal system utilization and application performance.

Important attributes of Cisco's HyperFlex portfolio highlighted are:

- **HyperFlex Compute.** HyperFlex supports the two most commonly deployed hypervisors: VMware vSphere and Microsoft Hyper-V. In addition, the Intersight Workload Engine (IWE) integrates a KVM-based hypervisor to give users the ability to run Kubernetes-based workloads without the need for third-party hypervisor licensing. HyperFlex clusters can be built with a mix of integrated appliances and with Cisco's broad portfolio of UCS servers. Users can scale their HyperFlex clusters linearly or choose to scale compute-only nodes without additional HCI licensing.
- **HyperFlex HX Data Platform.** Cisco's HX Data Platform provides native data services built from the ground up, specifically for HyperFlex. The HX Data Platform is an enterprise-grade, distributed file system designed with scalability, efficiency, application resiliency, data integrity, predictable performance, and high availability as core tenets. HyperFlex allows customers to incorporate external storage into their HX clusters, thus increasing flexibility and driving high levels of utilization rates for existing storage assets or leverage the HyperFlex cluster as scale-out storage for external VMs or containers through native iSCSI capability. The HX Data Platform provides the following services/capabilities: native data protection (with features like asynchronous remote replication included free of charge), the ability to provide RF3-level protection with as few as three nodes, a fully striped architecture that minimizes rebuilds, automatic data rebalancing, automatic self-healing recovery from node and drive failures, optional Logical Availability Zones (LAZs) to increase fault tolerance as you increase cluster size, nondisruptive rolling upgrades, zero overhead and instantaneous snapshots, built-in block checksums to protect against media errors, support for stretched clusters, and native encryption for security.
- **HyperFlex Networking.** HyperFlex offers a complete HCI platform that includes compute, HCI software, and fully integrated network in a single stack — all engineered to work together and supported by a single company. The key network functionality is provided by a fully integrated network, resulting in predictable performance, reduced complexity, reduced cost, low latency, and unified system management. Users can deploy HyperFlex clusters in conjunction with Cisco ACI for highly automated, software-defined, and policy-based networking capabilities. Users that deploy Cisco's ACI within their HyperFlex clusters will gain policy-based, application-centric management and orchestration of the HCI network fabric.

- **HyperFlex, an Intel Select Solution.** Cisco and Intel have partnered to deliver an Intel Select Solution for HyperFlex. This solution supports a range of performance-intensive, mission-critical applications on HCI, verified to meet or exceed a certain performance level.

Important attributes of Cisco's Intersight portfolio highlighted are:

- **Intersight Kubernetes Service (IKS)** is a lightweight container management platform for delivering multicloud, production-grade Kubernetes. IKS is aimed at simplifying provisioning, securing, scaling, and management of virtualized or bare metal Kubernetes clusters through the automation of processes that include integrating networking, load balancers, native dashboards, and storage provider interfaces. It has APIs for public cloud managed K8s offerings, including AWS Elastic Kubernetes Service (EKS), Azure Kubernetes Service (AKS), and Google Cloud Google Kubernetes Engine (GKE). IKS targets AI/ML development and data scientists looking for delivering GPU-enabled clusters and Kubeflow support via simple GUIs.
- **Intersight Virtual Appliance** enables a full-state backup of the data in the virtual appliance and stores it in a remote server. If there is a total site failure or other disaster recovery scenarios, the restore capability enables you to do a full-state system restore from the backed up system data.
- **Intersight Cloud Orchestrator** is an automation tool that enables IT operations teams to provide a consistent cloud-like experience for users. ICO brings public cloud and on-premises resources together with a solution that extends orchestration across any infrastructure and workload and integrates with the tools of choice. It saves time and streamlines automation with a GUI-based designer that simplifies the creation and execution of more complex workflows without the need of a coding expert. It also reduces risks by enforcing policy-based rules for what can be orchestrated and who can access workflows and tasks.

CHALLENGES AND OPPORTUNITIES

Decades of innovation have brought us datacenter solutions that are undeniably more capable than anything offered just a few short years ago. That said, too many datacenter teams continue to buy and manage their infrastructure the same way they did 10 or 20 years ago. IDC believes that this process has become untenable. IT departments must look to modernize operations by adopting best practices, and this includes HCI infrastructure tied to the edge and hybrid cloud services if they want to keep up with the unprecedented changes occurring all around them. Indeed, businesses of all sizes are looking to transform their companies to find new revenue streams, create deeper connections with their customers, or simply compete more effectively against new competitors that are unburdened by legacy systems and practices.

Technology suppliers that provide HCI systems must keep a pulse on requirements around compute, ephemeral and persistent data storage, and the kinds of workloads their customers run at their core and edge locations.

Meanwhile, I&O teams face several challenges when they deploy hybrid cloud management tools into production, such as silo-based teams using isolated open source tools, ineffective use of analytics, and poor visibility into application interdependencies across hybrid cloud environments. Additional challenges are related to IT culture, organizational structure, and technology.

Hybrid management offers several benefits for I&O teams to expand their business and strategic value and increase IT agility and speed. One of the major benefits is automation; specifically, automation of the configuration and deployment of containers and clusters and automation of capacity planning and workload balancing. Additional benefits are related to people, process, and technology.

A single source of truth via a single management plane provides critical information that all teams can use to identify and resolve problems across the complex array of physical, virtual, and container-based infrastructures, thus speeding time to problem identification and resolution.

More automated processes are enabled across container and infrastructure resources that utilize analytics to adapt infrastructure in real time and decide workload placement based on application requirements to drive operational simplicity and cross-silo/multicloud automation.

CONCLUSION

For many vendors, HCI is a building block for corporate hybrid cloud deployments. This is particularly important as business leaders depend on their IT departments to enable strategic initiatives designed to drive bottom-line improvements and better serve customers to generate new revenue streams that will ultimately give them a competitive edge in their industry. HCI and hybrid cloud have shown to have numerous advantages, including simplified life-cycle operations and streamlined single source support and enabling users to quickly integrate emerging technologies with predictable and repeatable results.

At the same time, HCI in support of hybrid cloud opens DevOps teams to cloud-based, SaaS data services in support of their containerized environments. Last, it is crucial to tie on-premises, HCI-cloud resources to hybrid cloud services through a single management plane in order to simplify monitoring and control planes.

MESSAGE FROM THE SPONSOR

Cisco and Intel have a long history of partnership, co-engineering proven data center solutions trusted by organizations around the globe. Today, that partnership continues to deliver optimized, modern hybrid cloud solutions. Cisco HyperFlex is powered by Intel Xeon Scalable processors, Intel Optane technologies, Intel Ethernet products to deliver consistent, predictable performance and reliability for critical on-prem workloads. Integrating advanced security features and Cisco Intersight management and cloud services, HyperFlex is certified as an Intel Select Solution for hybrid cloud use cases.

To find out more, go to cisco.com/go/hyperflex

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com