IDC

SPOTLIGHT

Sponsored by: Microsoft and Intel

Hyperconverged infrastructure is well suited to meet the needs of enterprises tasked with modernizing IT systems across core, edge, and hybrid cloud sites to achieve their digital transformation goals.

Hyperconverged Infrastructure Use Grows for Mission-Critical Workloads Across Core, Edge, and Hybrid Cloud Environments

June 2022

Written by: Carol Sliwa, Research Director, Infrastructure Systems, Platforms, and Technologies Group

Introduction

The use of hyperconverged infrastructure (HCI) combining compute, storage, and network virtualization technology is growing among small, midsize, and large enterprises. Recent IDC survey data shows that 74% of enterprises have moved one or more workloads from traditional SAN or NAS systems to HCI and close to 57% run mission-critical applications on hyperconverged systems. Many organizations replace traditional three-tier architectures with consolidated HCI as they refresh their aging IT infrastructure.

Enterprises increasingly take a hybrid, multicloud approach as they modernize their infrastructure in pursuit of digital transformation initiatives. Simplifying provisioning and management is especially important for digital-first organizations that need to stay agile. Streamlined HCI reduces complexity and offers flexibility to span on-premises core, edge, and public cloud sites. Software-defined HCI running on high-performance industry-standard server hardware can provide an economical platform for DevOps teams to test and ultimately deploy the container-based, cloud-native applications they build. IDC survey data shows that HCI has become a popular platform for running production workloads in containers. The technology can also help ease the migration of workloads and data between private and public cloud sites and offer a unified control plane to manage the two environments.

IDC research shows the most popular workloads currently running on HCI in core datacenter, colocation, or hosted environments include business intelligence/analytics, security applications, Internet of Things (IoT), databases, and business applications such as enterprise risk management, customer relationship management, and system control management. HCI use is also rising in edge environments, where the top use cases include IoT, business applications, security applications, and analytics.

AT A GLANCE

KEY STATS

According to IDC survey data:

- More than 70% of enterprises have moved workloads to HCI, and nearly 57% run mission-critical apps on it.
- Close to 70% of enterprises plan to shift toward HCI at edge sites.
- Only 1% of enterprises reported no HCI clusters connected to the public cloud.

KEY TAKEAWAYS

HCI use will grow as enterprises pursue digital transformation initiatives that require hybrid cloud infrastructure across on-premises, edge, and cloud sites.

Definitions

- **Hyperconverged infrastructure** is a software-defined system that combines storage, compute, and network virtualization technology within a single server under unified management. HCI software typically runs on standard general-purpose servers or prebuilt appliances, and it scales through the addition of nodes to a cluster. Some vendors support independent scaling of compute and storage resources.
- **Core** is the central datacenter containing an organization's primary storage and compute infrastructure. The core may be located at a company-owned site, a third-party colocation facility or, potentially, in the public cloud.
- **Edge** is an infrastructure site located away from the core datacenter in closer proximity to remote endpoint devices that generate or collect data.
- **Hybrid cloud** is an IT architectural approach that incorporates infrastructure located in one or more on-premises datacenters or colocation facilities and public cloud sites.

Benefits of HCI in Modern IT Environments

Converging compute, storage, and network resources into a single HCI system can reduce complexity and costs for enterprises accustomed to three-tier architectures that often need to be provisioned, operated, and maintained by IT staff with specialized skills. Simplifying deployment and management can be especially advantageous for organizations facing IT staffing challenges at edge, branch, or remote office sites that require no more than a few servers to meet their infrastructure needs. Software-defined HCI running on industry-standard hardware also enables users to scale out quickly and cost effectively, or burst to the public cloud, whenever their needs change. HCI systems increasingly integrate with public cloud platforms and give customers the ability to link or manage core, edge, and cloud resources through a unified, user-friendly console to ease their transition to a hybrid cloud strategy.

Streamlined HCI systems can be advantageous for data processing and analysis in edge environments located in closer proximity to IoT and endpoint devices than core datacenters. They have also proven beneficial in easing large workload consolidation and migration projects that may be difficult to size and in simplifying virtual desktop infrastructure (VDI) that has grown complicated to scale. HCI can help organizations take advantage of the latest compute, storage, and networking technology advancements to meet the performance needs of demanding workloads such as databases and analytics.

Key Trends

IDC research shows that HCI adoption is broad among small, midsize, and large enterprises. Most organizations deploy HCI in core datacenters and colocation facilities today, but survey data shows that edge use is common and likely to grow, especially for workloads such as IoT data processing, business applications, and analytics. The most important attributes organizations consider when deploying HCI at the edge include security, integration with public cloud services, combined data protection, and systems management automation.

Software-defined HCI often serves as an on-ramp to the public cloud. IDC survey data showed that 99% of respondents connected an HCI cluster to a public cloud. More than 60% run separate production workloads on HCI and public clouds connected through a unified management console. Many organizations indicated they access public cloud resources for short-term bursting or seasonal needs, application development and testing, disaster recovery, backup, and long-term archiving.

Microsoft Turns Focus to Cloud-Managed Azure Stack HCI

Microsoft's continued growth from personal computing into enterprise software and public cloud services has given the vendor a sizable technology footprint in many IT organizations. According to IDC research, Microsoft's public cloud infrastructure-as-a-service (IaaS) revenue soared 45.1% from 2019 to 2020, accounting for $9.4 billion of the $65.5 billion worldwide public cloud IaaS market. Microsoft's enterprise application revenue grew 10.7% from 2019 to 2020 to reach $5.3 billion. During the same time frame, Microsoft's revenue in the worldwide software-defined storage controller software market that includes HCI increased by 21.6% to reach $243.3 million.

Microsoft initially offered Azure Stack HCI through a datacenter software feature with its 2018 introduction of Windows Server 2019. However, the company soon pivoted to a cloud service–based HCI model designed to enable customers to centrally manage the product through an Azure portal and connect to Azure cloud services. Although Microsoft continues to support the Windows Server 2019-based HCI option, it now focuses on the redesigned Azure Stack HCI that uses the same brand name. Organizations that deployed the initial Windows Server-based HCI will need to do a fresh installation of the new Azure Stack HCI software if they want to take advantage of its differentiating features, including enhanced security and cloud services such as Azure Kubernetes Service (AKS) and Azure Virtual Desktop that they can use on premises.

Released in 2020, the newer HCI software is part of Microsoft's Azure edge infrastructure portfolio and Azure Arc technologies designed to bring cloud services and security to hybrid and multicloud environments. Although customers manage and monitor the redesigned Azure Stack HCI through a central cloud control plane, they run the software on premises on industry-standard server hardware. Microsoft partners with OEMs to offer more than 425 validated hardware options.

Since Microsoft delivers the redesigned Azure Stack HCI as a service, as opposed to a separately licensed software product, customers need an Azure subscription to access billing, support, and management through a central Azure portal. Customers use the same wizards to provision and register virtual machines (VMs) for an Azure Stack HCI-based on-premises deployment that they use to launch VMs to an Azure cloud region, and they can apply the same security and governance to VMs in spanning on-premises and public cloud sites through the cloud management console. They also gain immediate access to security updates and new features since Microsoft delivers Azure Stack HCI as a cloud service.

Customers most commonly use Azure Stack HCI in two-, four-, or eight-node configurations, and Microsoft offers options to federate large installations into a single pool of resources. Customers in regulated industries, such as financial services and healthcare, often require on-premises infrastructure, and users running high-performance workloads generally need to keep data close to the systems that process it to minimize latency. One of the most popular use cases is low-cost edge infrastructure at remote retail and manufacturing sites that, in some cases, require only two-node clusters for high availability or run container-based applications that make use of AKS. Other top workloads include Azure Virtual Desktop deployments, SQL databases, and the "lift and shift" of Windows Server 2008 and 2012 Hyper V-based workloads to new infrastructure. Moving to Azure Stack HCI comes with the added benefit of free extended security updates on old and unsupported operating systems.

Intel Is Significant HCI Partner for Microsoft

Microsoft made performance optimization a key area of focus in its software development work, and key partners offer advanced hardware options. More than 75% of the Azure Stack HCI validated nodes and pre-integrated system choices use Intel technology, and the CPU market leader's latest third-generation Xeon Scalable processors boost performance over prior alternatives. Sixth-generation Intel Virtualization Technology provides an abstraction layer to ease workload migration between HCI nodes or clusters with differing generations of processors. Customers also stand to benefit from the Intel Select Solutions program, in which Intel works with Microsoft and OEMs to test and benchmark Azure Stack HCI hardware to ensure it is optimized and ready to deploy.

HCI customers with demanding performance- and latency-sensitive workloads may want to consider Intel's network adapters and Optane storage options. Intel's latest 800 Series network adapters support remote direct memory access (RDMA) to accelerate inter- and intra-cluster communication. The Optane products use nonvolatile 3D XPoint memory technology that Intel codeveloped with Micron to bridge the performance and price gap between faster, more costly volatile DRAM and slower, less expensive nonvolatile NAND flash. Intel Optane is available in two form factors: standard solid state drives (SSDs), which connect via the PCI Express storage bus, and dual inline memory modules (DIMMs), which attach directly to the memory channel. The Optane SSDs promise more consistent high performance than flash-based SSDs and target write-intensive workloads such as online transaction processing, high-performance computing, data caching, and logging. Systems with Optane SSDs generally use the drives as a cache tier and have higher-capacity NAND flash-based SSDs for data storage.

The Intel Optane Persistent Memory (PMem) DIMMs operate in one of two ways: Memory Mode, to expand the available pool of volatile system memory at lower cost than DRAM, or App Direct Mode, to enable applications to access a tier of nonvolatile memory for data storage at lower latency than Optane and flash-based SSDs that require the data to move to and from the storage bus. Applications require no changes to use Optane PMem in Memory Mode, but Intel had to work with partners and providers on system and application modifications to facilitate App Direct Mode. Workloads that can benefit from Optane PMem DIMMs in App Direct Mode on systems such as HCI include in-memory databases, analytics, and artificial intelligence.

On the security front, Intel offers several technologies with relevance to HCI installations. Intel Crypto Acceleration can help speed encryption and minimize the performance impact. Intel Multi-Key Total Memory Encryption uses the AES XTS algorithm with 128- or 256-bit keys to encrypt all memory that the processors access, and key management software enables granular encryption at a page or VM level. Intel also joined other Microsoft partners to do the necessary integration work to support Microsoft's Secured-core server initiative to protect HCI infrastructure against firmware-level attacks through hardware-based root-of-trust technology.

Challenges

IDC survey data shows that IT leaders who have not replaced SAN and NAS systems with hyperconverged infrastructure harbor fears that HCI lacks capacity and performance density on a per-node basis, does not independently scale compute and storage resources, and fails to meet their security requirements. HCI users noted scaling challenges to maintain a single view of all clusters in a multivendor environment, to guarantee application performance for certain workloads, and to achieve sufficient storage performance without excessive cost. Close to 60% of respondents indicated their HCI systems struggle to maintain consistent performance with large data sets, especially between cluster nodes, and they face high network bandwidth costs to keep data consistent between core and edge HCI platforms.

Microsoft is a relatively new entrant in a market that includes well-established players with mature products that link to competitors' clouds, but its broad range of software and hardware options, including the performance-focused technology of key partner Intel, stands to address many of the concerns that IT organizations raise about HCI.

Conclusion

The use of HCI will grow as enterprises increasingly take a hybrid cloud approach with digital transformation initiatives that require modern infrastructure designed to span on-premises datacenters, edge sites, and multiple public clouds. Flexible, software-defined HCI systems can simplify provisioning and maintenance, connect to public cloud resources and services, and offer a unified control plane to manage infrastructure that spans private and public cloud sites. They also deploy on economical industry-standard hardware to offer a cost-effective platform to test and deploy the container-based, cloud-native applications that DevOps teams in digital-first organizations increasingly build.

The use of HCI will grow as enterprises increasingly take a hybrid cloud approach with digital transformation initiatives.

HCI users have identified concerns in guaranteeing consistent performance, scaling storage and compute resources, and maintaining security at an affordable cost. Vendors that concentrate on solving the challenges that HCI users confront have a significant opportunity for success. With Azure Stack HCI, Microsoft and Intel have prioritized features and technology that can help address customer concerns.

About the Analyst

Carol Sliwa, *Research Director, Infrastructure Systems, Platforms, and Technologies Group*

Carol Sliwa is a Research Director for Storage Systems in IDC's Enterprise Infrastructure Practice. Her core research area spans block, file, and object storage and hyperconverged infrastructure. With more than 25 years of experience as a technology journalist, including 13 years covering enterprise storage, Carol gained extensive insight into the ways in which the industry has adapted systems over time to address the evolving needs of IT customers.

MESSAGE FROM THE SPONSOR

More About Microsoft and Intel

In partnership, Microsoft and Intel make it easy to support the evolving needs of businesses with compatibility, security, and flexibility for their infrastructure. Azure Stack HCI is one of the offerings coming from this partnership. Learn more at www.azure.com/hci

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com