

Hyperconverged infrastructure (HCI) brought significant simplicity to infrastructure deployment and management. Now, IT managers want to leverage the HCI experience across more workloads. New requirements in these workloads are driving an evolution in the definition of HCI.

Expanding the HCI Experience to More Workloads: Factors to Consider

August 2020

Questions posed by: HPE

Answers by: Eric Burgener, Research Vice President, Infrastructure Systems, Platforms, and Technologies

Q. What is the importance of IT agility to enterprises undergoing digital transformation (DX)?

A. Digital transformation is the evolution of enterprises of all sizes toward more data-centric business models that leverage big data and analytics to improve internal workflows, products, and services as well as drive disruptive improvements in the ability to compete and open up new markets. Business in the digital ecosphere is more dynamic and moves faster than business in the older, more static era. To compete, IT organizations are implementing new methods such as DevOps, leveraging new technologies such as artificial intelligence and machine learning (AI/ML), and deploying new approaches such as the Internet of Things (IoT). This evolution requires agility and the ability to respond rapidly to changes in demand and newly identified market opportunities. IT has become more important than ever in providing the enabling infrastructure to support DX.

The need for better agility translates to very specific requirements for IT infrastructure and operations. Provisioning, performance, availability, scalability, configurability, and functionality are all impacted, and enterprises require new storage architectures and technologies to meet these requirements. Agility has also impacted consumption models, and customers are looking for more flexible ways to acquire infrastructure that include subscription-based, pay-as-you-go approaches.

Q. What is the hyperconverged infrastructure (HCI) experience, and what is driving such interest in it?

A. Over the past decade, HCI revenues have grown rapidly. In many cases, IT organizations are replacing legacy infrastructure with HCI during technology refreshes. The simplicity of traditional HCI represents an experience in contrast to the more complex model that has been associated with deploying traditional three-tier enterprise infrastructure in the following areas:

- » **Purchase and deployment.** The HCI experience is extremely simple: A working system composed of servers, storage, networking, and infrastructure software is ordered under a single SKU from a single vendor with all the components prevalidated to operate together. Installation is often automated through templates provided by the vendor, improving the speed and reliability of deployment. In addition, HCI allows customers to start small and grow easily, providing more flexibility for its use in edge and core deployments.
- » **Simple, unified management.** Traditional HCI offers a unified management interface where the compute, storage, and networking resources are managed from a single unified interface. The software-defined nature of HCI makes it more flexible, and the fact that all components come from a single vendor enables ease-of-use features such as automated configuration, remote deployment and management, virtual machine (VM)-centric provisioning and data services application, and one-click software upgrades — all of which are useful in cases where sophisticated storage management expertise is not available.
- » **Single point of support.** When it comes to HCI configuration, a single point of support means that customers do not need to deal with separate vendors for compute, storage, and networking resources. It eliminates the finger-pointing that can occur in infrastructure composed of different products from different vendors.

These aspects of the HCI experience drive real value for enterprises, which now want to apply HCI to different types of workloads in their datacenters. Traditional HCI is a good fit for environments where workload growth is more predictable, while infrastructure that enables independent scaling of resources and supports higher-performance, availability, and efficiency objectives can be a better fit for a wider class of workloads that are more unpredictable in terms of growth.

Q. How is the definition of HCI evolving, and why?

A. IT managers are looking to apply the HCI experience to more workloads. The definition of HCI has been evolving, however, particularly to address perceived challenges with the traditional HCI model around certain performance, availability, and agility issues. While enterprises find the simplicity of the HCI model very attractive, they see its traditional performance, availability, resiliency, and efficiency characteristics as negatives for certain types of common enterprise workloads.

Probably the biggest concern with traditional HCI's ability to handle larger enterprise workloads, particularly those with unpredictable growth paths, is that compute and storage resources can be added only in fixed ratios with the addition of a new HCI node. There is little flexibility to add either resource independently, resulting in a situation where one of the resources is always overprovisioned relative to the other. This imbalance can begin to add significant cost as configurations scale in size (in terms of node count and overall capacity). As customers expand the size of their HCI configurations to accommodate a broader set of workloads, this lack of configuration agility is a concern.

Maintaining consistent performance can also be a challenge when an application's data set is too large to fit on a single HCI node. When data outside of the node on which an application is running must be accessed, network latencies are incurred that can add noticeably to perceived application response times, particularly when using all-flash configurations. Additional network latencies may be incurred even for data sets that might fit on a single HCI node if that data is spread across multiple nodes to meet data resiliency objectives. Dual-parity RAID schemes that maintain data resiliency even in the face of two simultaneous storage device failures are a baseline requirement for many enterprise workloads that customers would like to move to HCI, but spreading data across multiple nodes to achieve that level of resiliency forces traditional HCI customers to incur additional network latencies. Achieving sufficient levels of resiliency may also require a minimum HCI node count that is higher than the minimum HCI node count required purely to service the application workload. Recovering quickly from node failures can also be time consuming, especially when data is spread across nodes (whether for capacity or resiliency reasons).

One solution to these challenges is a "disaggregated" HCI model that pairs the HCI experience with disaggregated compute and storage resources. In this model, a vendor provides a unified management interface that lets customers centrally manage all resources in a VM-centric manner from a single plane of glass, includes integration and automation that simplify deployment and ongoing management, and offers a single point of support. This configuration lets customers scale compute and storage resources independently, using a proven enterprise-class storage array with comprehensive storage functionality, support for "five-nines plus" availability, and the efficiency of shared storage. By addressing the independent resource scaling and enterprise-class functionality issues, disaggregated HCI solutions can host the more demanding workloads that are not a good fit for traditional HCI with higher resource efficiency.

Q. What is VMware Virtual Volumes, and why is this feature important?

A. VMware Virtual Volumes (VVols) is a management API in VMware vSphere that effectively enables VM-level management within block-based arrays. Without this API, storage capacity in external arrays must be managed as volumes, making it difficult to perform storage monitoring, troubleshooting, and management tasks at the VM or application level. With VVols, array monitoring can be done at the VM level and high-performance, array-based data services can be easily applied on a per-application basis even in cases where multiple VMs are hosted on a single volume. VVols significantly improves the performance, manageability, and efficiency of resource utilization for virtual environments built around VMware's vSphere technology that are running on block-based external storage arrays.

VVols is an API, however, and array vendors have to make modifications to their block-based array management tools to enable its use. Vendor implementations can vary significantly in their functionality. If VM-level management when using array-based data services is important when evaluating a vendor storage platform, enterprises will need to ensure they understand the comprehensiveness of each vendor's VVols support. Features to look for include application provisioning

templates and policies, quality of service, snapshot portfolio support (space-efficient snapshots and snapshot backup integrations), thin provisioning, encryption, deduplication, automatic space reclamation, rapid restore options, persistent tiering, a replication portfolio (including VMware Site Recovery Manager integration), and public cloud integration/support. The completeness of a storage array vendor's VVols integration can significantly impact performance, efficiency, productivity, and cost.

About the Analyst



Eric Burgener, Research Vice President, Infrastructure Systems, Platforms, and Technologies

Eric Burgener is a Research Vice President within IDC's Infrastructure Systems, Platforms, and Technologies Group. Mr. Burgener's core research coverage includes storage systems, software and solutions, quarterly trackers, and end-user research as well as advisory services and consulting programs. Mr. Burgener's research includes a particular emphasis on emerging infrastructure technologies, including all-flash arrays (AFAs), persistent and storage class memory, and software-defined storage. He is an active participant in the IT Buyer's Research Program at IDC and blogs throughout the year on the topic of infrastructure and data management.

MESSAGE FROM THE SPONSOR

HPE Nimble Storage dHCI is a disaggregated HCI platform that delivers the HCI experience, but with more flexibility and better resiliency, performance and efficiency than traditional HCI. The platform simplifies infrastructure from install to upgrade with customer self-installation in 15 minutes from rack-to-apps, VM-centric operations through VMware VVols, and one-click software upgrades. It also brings the flexibility to scale resources independently to maximize efficiency, provide 99.9999% of data availability and sub-millisecond latency for business-critical applications and mixed workloads, and lower total cost of ownership.

To maximize IT agility, HPE Nimble Storage dHCI goes beyond traditional HCI with the industry's most advanced AI for infrastructure, HPE InfoSight, and a cloud consumption experience through HPE GreenLake. HPE InfoSight uniquely provides HPE Nimble Storage dHCI customers comprehensive insights into their virtual infrastructure that predicts and prevents problems, diagnoses performance issues, forecasts resource needs, and optimizes resource utilization. Combined with HPE GreenLake, HPE Nimble Storage dHCI delivers an intelligent foundation for private cloud enabling VMs as a service.

The logo for IDC Custom Solutions, featuring a blue circular icon with a white dot in the center, followed by the text "IDC Custom Solutions" in a sans-serif font.**IDC Research, Inc.**

5 Speen Street
Framingham, MA 01701, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2020 IDC. Reproduction without written permission is completely forbidden.