



Global Knowledge®

Expert Reference Series of White Papers

How Hyper-
Converged
Infrastructure Can
Change the Way
You Deploy IT
Systems

How Hyper-Converged Infrastructure Can Change the Way You Deploy IT Systems

Brian Eiler, VCI, VCIX, PCNSE, ASE, CCNA

Introduction

What if I told you that you could squeeze a data center into a box the size of a small rack-mountable server? Imagine what you could do with it. You could store it in a closet, ship it to branch offices, or even buy an extra one and send it away for disaster recovery.

But could one little box really hold enough power to run an entire data center? Maybe, but these aren't magic boxes. They are, however, designed to work like building blocks. As you grow, you add more boxes. But unlike traditional servers, as you add these boxes, you are effectively scaling out a complete data center with compute, storage, and networking power. Your capacity and performance are actually expanding together!

If all this sounds like pure magic or simply buzzwords, keep reading because you're in for a surprise. This is the very real world of Hyper-Converged Infrastructure (HCI), and it will change the way you deploy IT systems.

The Challenge

If you have ever deployed equipment in a data center, you're already familiar with some of the challenges. For everyone else, let's review why adding equipment into data centers can go from a nice step-by-step list on a whiteboard to a frustrating experience with your hands stuck in a rack at two in the morning.

Traditional Systems

First, let's discuss traditional systems in the data center: the racks of individual servers, storage arrays, and network equipment. These are discrete components typically sourced from multiple vendors or manufacturers. The server is made by one company, the storage by another, and the networking by a third company. The combinations of vendors and products seems nearly endless, which is what makes this model so flexible and desirable to many people.

This flexibility, however, is just an illusion. Realistically, due to nuances in the implementation or subtle differences in design, there are far fewer working combinations of these components than one would expect. The vendors themselves realized this problem and provide Hardware Compatibility Lists (HCL) to help customers determine which components work together. Unfortunately, not everyone reads these lists, and even those who do read them may make incorrect assumptions about seemingly innocuous things like firmware versions that later prove to be disastrously incompatible.

Even if you have a winning combination of products between your servers, storage, and networking, you will still face the challenge of performance bottlenecks. Even new systems may have a bottleneck on one of the resources, it just may be insignificant at the moment. As time passes, the bottlenecks become more pronounced. Yet in traditional systems, no single management tool exists to help diagnose or identify which component is causing the bottleneck. Troubleshooting turns into a game of chasing the weakest link in the resource chain until you finally upgrade nearly every subsystem.

Converged Systems

To combat challenges in traditional systems, vendors began creating what are known as converged systems. These started as server-blade chassis and enclosures that include servers and networking equipment. Some later versions even include a pool of storage that can be assigned to servers. One of the best-known forms of convergence is via storage networking, specifically a storage protocol known as Fibre Channel over Ethernet (FCoE).

FCoE systems leverage special converged networking and storage adapters to allow servers to access Fibre Channel storage devices using a single set of Ethernet cables. The chief advantage to these converged systems is the ability to use very thin servers or blades to physically accommodate the number of hardware cards required to access external storage while still maintaining high speed access to the network.

While converged systems help reduce some data center cabling, they don't necessarily reduce the complexity of managing the data center overall. With converged systems, it is still necessary to manage servers, storage, and networking as separate components even if the connectivity between them has been consolidated.

Hyper-Converged Infrastructures

HCI systems are rack-mounted devices that are typically no larger than a 1U or 2U rack-mounted server (approximately 3.5 inches high in a standard 19-inch rack). They look like small servers, yet they also bear a resemblance to a miniature blade chassis. These devices often have many hot-pluggable modules in the front and back for things like hard drives and server blades. Depending on the make and model of the HCI system, the device may only be tethered to your data center by four wires: two power cables and two Ethernet cables.

HCI represents the third generation of systems intended to simplify data center management. In the next section of this paper, we will explore how HCI systems deliver on that promise.

The Advantages of a Hyper-Converged Infrastructure

Taking away marketing terminology, at the most basic level an HCI system is simply a machine containing all the compute, storage, and networking resources to host a set of virtual machines. It's natural to wonder how this differs from every other server out there with local storage and a hypervisor. While an HCI system may look like a typical rack-mounted server, you won't manage it in the same way. HCI systems vary by vendor, but inside, you will typically find one or more small form-factor blade servers, a shared storage system, a central management module, and an Ethernet switch that connects everything inside together and then exposes the servers to the outside world. Finally, these server blades boot to a hypervisor that hosts virtual machines stored on the shared storage system thereby creating a virtual data center in a box.

Part of that design may sound familiar. Vendors have been producing blade chassis units that provide similar functions for over a decade. However, blade server chassis are unable to solve the bigger picture problems of operational management. For that, we have HCI. Here are a few areas that have helped propel HCI to the forefront of the battle for the data center:

- Management on a single pane of glass
- A building-block approach with linear scalability
- A single-vendor stack of equipment
- Managing the application performance experience

Management on a Single Pane of Glass

The old saying "knowledge is power" may be true, but knowledge can only be powerful if you can find it and access it. For instance, in traditional systems we can see an incredible amount of detail about our servers, storage,

and networking, but you typically must open several different management consoles to get the information. There is also the additional problem of aggregating information and correlating events between the different silos of hardware.

As previously mentioned, HCI systems include the compute, storage, and network resources within their chassis, and all of these resources are controlled by management modules within the HCI systems. This allows the HCI system vendor to provide administrators with a unified management interface that shows various resources, their current utilization, and potentially their expected capacity. This view has become known as the “single pane of glass” because you can manage all aspects of the virtual data center from a single screen.

A Building Block Approach with Linear Scalability

HCI systems are designed to be the building blocks of your data center. Each chassis is a self-contained miniature data center with a set quantity of CPUs, RAM, storage, and network connections that can be provisioned to virtual workloads hosted and managed by HCI system software. Though the HCI systems have all the resources to allow them to stand on their own, they really shine when linked together.

Unlike traditional systems, expanding an HCI solution is an incremental process of installing an additional chassis as opposed to buying trays of storage, stacks of switches, or racks of servers. You can then configure the HCI management software to link the chassis together, expanding the hardware resource pools in unison. At the end of the expansion, the additional compute, storage, and networking resources in the new chassis become available to the virtual machines.

Most HCI vendors also allow virtual machines to leverage storage in different chassis thereby permitting a virtual machine to run in one chassis while its storage may exist in another. Consider a case where the CPU and RAM were fully utilized in one chassis, yet it had storage capacity that wasn't fully used. This process helps ensure the hardware in each chassis can be fully utilized. Some vendors have taken this approach even further by striping the storage across multiple chassis, resulting in significant performance gains each time you expand the HCI system because more disk devices are handling the storage I/O.

A Single-Vendor Stack of Equipment

Unlike traditional or even converged systems where the servers are typically provided by a different hardware vendor than the storage or the networking, HCI systems are built and delivered by a single vendor. In some cases, the hypervisor and management software will even be provided by that same vendor. The disadvantage to this approach is the potential lock-in that you might feel as you continue building out your HCI solution. If you wish to expand the existing cluster and continue taking advantage of that single pane of glass management, you'll need to continue purchasing hardware from the same vendor. This disadvantage is in some ways mitigated by the building block approach of HCI. You won't give up linear scalability by building a second or third silo of HCI systems, but you would likely need to manage each vendor's solution separately. Even then, separate vendor silos could be a desirable outcome if you need to keep an air-gap between different projects or internal customers within your organization. Each silo would be a self-contained mini data center that can scale separately with that project or customer.

From an operations perspective, a single-vendor solution becomes extremely advantageous during troubleshooting situations because you have a single company to call for help. Plus, you are unlikely to encounter scenarios where tech support tells you to call another company because they don't support part of your hardware configuration. In HCI systems, when a single vendor designs, builds, and supports all the hardware and software of the system, you aren't likely to add components the designers didn't anticipate. Remember, these are self-contained systems. HCI systems are typically not attached to external silos of storage nor are attempts made to add special cards to these servers. To add capacity to your mini data center, you will add more HCI chassis and link them together over the network.

Ultimately, the single-vendor design may be enough to drive some organizations to use HCI because it all but removes the vendor to vendor finger pointing when things don't work. That alone can result in a reduction of downtime.

Managing the Application Performance Experience

Perhaps one of the most appealing, yet least advertised, features of HCI systems is the ability to manage the application performance experience. Using policies that stipulate the performance requirements of virtual machines, administrators can provision applications with a sense of certainty that the performance delivered will match the performance promised.

This approach may sound familiar because the concepts come from the Software Defined Data Center (SDDC). These policies are based on the different types of hardware provided in the HCI systems, and the policies are often grouped into tiers much like the way public cloud providers sell servers with different levels of performance or availability. When you assign a policy to a virtual machine, the requirements specified in that policy will be used by the HCI system to determine where the virtual machine will execute and store its data. Later, if the virtual machine must be moved, the new hardware is first checked against the policy for compliance to ensure the application performance will not be compromised.

HCI and the Software Defined Data Center

HCI really is the Software Defined Data Center (SDDC) in a box. HCI systems leverage the innovations that make up the SDDC: Virtualization, Software Defined Networking (SDN), Software Defined Storage (SDS), and Policy-Based Management (PBM). In this section, we explore how HCI leverages each innovation to become the true building blocks of the SDDC.

Virtualization

Server virtualization is at the heart of HCI architecture. A piece of software that acts as both an operating system and an application is installed on each server node within the HCI chassis. The hypervisor is responsible for hosting the virtual machines and providing them hardware access such as CPU, RAM, storage, and networking. This hardware abstraction makes the virtual machines ignorant to the make and model of the underlying physical hardware. This, in turn, makes it possible for the virtual machine to move to different hardware even while the applications are running.

The HCI management software controls the hypervisors within each chassis and allocates the hardware to the virtual machines based on the policies defined by the systems administrators. There are several different hypervisors on the market, each with their own strengths and weaknesses. Some HCI vendors rely on commercially available hypervisors such as VMware's ESXi or Microsoft's Hyper-V whereas others leverage open source hypervisors such as Linux KVM. There are even HCI vendors that have created their own hypervisor, such as Nutanix Acropolis.

While choosing an HCI vendor based on the hypervisor may seem logical and feel like an important decision, keep in mind that HCI systems are typically managed using a proprietary user interface instead of the hypervisor's traditional management interface. Depending on the automation and policy features in the HCI solution, you may decide to keep it separate from your existing virtual infrastructure.

Software Defined Networking (SDN)

SDN and its close cousin Network Function Virtualization (NFV) have made it possible for organizations to replace their physical network devices like routers and firewalls with equivalent virtual network devices deployed as virtual machines on top of hypervisors. These virtual routers and firewalls can be quickly deployed and configured through automation routines in cloud management platforms, so that as an application is deployed all the security and connectivity components are deployed with it.

HCI solutions can be deployed with virtual firewalls, routers, and VPN devices that replace some of the traditional hardware found in both remote offices and data centers. The added benefit to using HCI with virtual network devices is the automatic, high availability of these virtual devices. Given the level of hardware redundancy present in HCI systems, you are automatically protected from many common hardware failure problems. Additionally, you can deploy multiple virtual devices if you wish to mitigate potential software or configuration related outages.

Software Defined Storage (SDS)

Of the SDDC components that have propelled HCI to the front lines, Software Defined Storage (SDS) has had the largest impact. SDS is an innovation in two areas. First, it changes the way that storage is written to disks. Second, it changes the way servers consume storage.

Though SDS can achieve storage virtualization and abstraction using a variety of methods, the method used in HCI systems is the virtualization of locally attached storage devices in servers. This is a large departure from the traditional data center model where storage is centralized to a purpose-built piece of hardware called a storage array or Storage Area Network (SAN). Instead of relying upon hardware-based erasure encoding at the array level, such as RAID-5 or RAID-6, SDS controls redundancy and performance through policy. These policies determine how many copies of a data block will exist and where that block will be written. SDS also enables data blocks to be written across multiple servers in the data center thereby optimizing space utilization and increasing disk performance. Most SDS systems leverage part of each server's RAM as a cache as well as solid state disks to further increase performance.

The other side of SDS involves changing how servers consume the storage. In a traditional system, storage devices are formatted and presented to hypervisors as block storage or network storage devices with file systems. In either case, the servers view this storage as a list of files and folders. SDS changes this approach by presenting the servers with a flat volume of objects. Each object represents a part of a virtual machine, including the virtual disks, configuration files, and snapshots. The administrator assigns storage policies to the virtual machine, which the hypervisor then uses to determine where the object's blocks should be stored. The location is not based on a file system but rather an abstract grid of storage locations on different servers and storage devices. This abstraction of storage is transparent to the administrator and allows the HCI management software to move storage blocks as needed to satisfy the requirements in the storage policies.

Policy Based Management (PBM)

Whether you operate a small, medium, or large data center, you have customers and their satisfaction defines your success even if those customers are your fellow employees. Delivering what your customer expects (or demands) is hard enough, but in the data center world, you must deliver it constantly. Unexpected performance changes impact many people: your customer, their customer, perhaps your other customers, and, of course, you. After all, you'll be the one getting the phone call in the middle of the night. So how can HCI help you avoid performance surprises? Let's look at an example.

Imagine that you are the administrator of a medium-sized data center, and you have been asked to create a new virtual machine for an application server. The request includes a detailed set of hardware requirements like the number of CPUs and the amount of memory and hard drive space. Additionally, the requestor has noted that this application is extremely sensitive and must be placed on the highest performing equipment possible.

Being the data center administrator, you have a fairly good idea which servers, storage, and networking devices perform the best in your data center. You work with these systems every day and have come to realize that certain disks in specific arrays outperform others. You also happen to know which hypervisor hosts are the least busy. Taking this knowledge into account, you quickly provision the virtual machine on the fastest server, storage, and networking devices available. Mission accomplished. Or was it?

Sure, you placed the virtual machine on the best hardware available, but fast-forward a week or a month. Is that hardware still the best performing equipment, or have others had the same thought you did and thereby overburdened that high-performing equipment?

Let's take a different approach to this example. You are still the administrator of the data center, but instead of having years of experience, you will be a new-hire with little to no familiarity with the past performance or the underlying design of the hardware. This time we're at the mercy of the site documentation (assuming it exists) and whatever performance charts might be available.

What we need is a system that can help us choose the right hardware resources for the job by identifying the hardware that matches the provisioning requirements of our servers. Ideally, this system also prevents accidentally moving the server to non-compliant hardware in the future. To further protect the performance of everything using this hardware, the system must also prevent unauthorized workloads from consuming these precious hardware resources.

Fortunately, there is such a system, and it is known as Policy Based Management (PBM). Though each vendor may implement it in slightly different ways, PBM remains a core aspect of HCI systems. PBM enables the HCI systems to manage and optimize the use of all the hardware resources within each chassis as well as across all the interconnected chassis. PBM ensures that the right virtual machines have access to the right resources instead of leaving pools of unused disk space or idle CPUs around the data center. Optimizing the hardware across all the interconnected chassis allows HCI systems to achieve true linear scalability as you add chassis to the system. PBM is, therefore, one of the most important aspects of HCI and is a major advantage over the traditional server/storage silo approach.

Policy Based Management in Action

To help illustrate how PBM is leveraged in HCI systems, let's consider the following example.

An administrator acquires a new HCI system and then connects to its management interface to review the inventory of hardware resources. Using the HCI management software, the administrator reviews the identifying labels called "tags" assigned to hardware resources within the HCI system. Each tag represents a specific identifying characteristic or capability that the administrator could use later to filter the hardware list. For example, the tag "solid-state-disk" could be used to identify high performance storage. The tag "replicated-disk" could identify storage that is being automatically replicated to a disaster recovery location. Typically, HCI systems are pre-programmed by the vendor with all the necessary tags assigned to their hardware.

Next, the administrator creates one or more policies that will later be assigned to the virtual machines. These policies control which pieces of hardware will be available to virtual machines and may include several tags for different capabilities. For instance, a policy named "Fast and Mirrored Storage" with both the "solid-state-disk" and "replicated-disk" tags would only permit the virtual machine to use hardware possessing both tags. Finally, the administrator makes these policies available to the people who create virtual machines. Many HCI systems come preconfigured with policies to simplify administration.

Later, when another administrator creates a new virtual machine, he or she assigns one or more policies to the virtual machine based on the application requirements. The policy assignment works as a filter on the hardware list, ensuring that the administrator can only select hardware that meets the policy requirements assigned to the virtual machine. In this way, PBM provides a type of guarantee that the right resource will be assigned to the right virtual machine.

Furthermore, this policy is bound to the virtual machine and remains in effect even if moved throughout the HCI system. This ensures the virtual machine will remain on the correct type of hardware even if it is moved to another HCI chassis.

What Does HCI Mean to Me and My Organization?

When is the right time to make the move to HCI? Unfortunately, there isn't a clear-cut answer that applies to everyone. There isn't a specific audience or type of customer where HCI would be best suited, but there are some indicators that you have a valid use case for HCI.

First, if your data center has many partially-filled silos of storage and servers, you're likely not utilizing the storage or servers to their fullest potential. This is a common occurrence, especially in storage systems, because storage is carved into volumes that are presented to servers. These volumes are blocks of disk-space written to specific disks in an array. Unlike the automatic block movement found in SDS systems, data blocks aren't easily moved in traditional storage arrays.

You may also have clusters of servers that are either over- or underutilized because the servers have no access to all the storage in the environment. In HCI systems, these hot spots are avoided by presenting all the hardware to the management system, which then allows the administrator to properly balance the loads or create policies to automate the process.

Another area to consider involves people. Think about the skillset of your team. HCI systems are designed to remove the complexities of managing hardware. However, if you have a redundant staff of dedicated personnel that exclusively manage server hardware, storage arrays, and networking hardware, then the traditional model may not be a hindrance to the agility of your IT department. The typical IT department doesn't fit that picture because employing a truly redundant workforce is expensive and hard to maintain.

Finally, perhaps you have an upcoming project that should remain isolated from the main data center or you have remote locations where it would be appropriate to ship them a data center in a box. Because HCI systems can often be maintained entirely separate from your existing environment, these systems are well-suited for drop-in deployments. Plus, you can always scale them out as the project or office grows.

HCI Vendors

HCI represents a major change to the way we deploy virtual infrastructures and the applications within those infrastructures. As you might expect, the HCI market is filled with vendors, each with their own strengths, weaknesses, and approach to HCI. These vendors, products, and solutions change rapidly, so be sure to evaluate several products before making your final decision. Below is a list of a few vendors to get you started in your research, but keep in mind this list it is not intended to be exhaustive.

- [Atlantis Computing](#)
- [Cisco: Hyperflex](#)
- [Cisco/IBM: VersaStack](#)
- [Dell/EMC: VxRAIL](#)
- [VMware](#)
- [Diamanti](#)
- [HPE/Simplivity](#)
- [HyperGrid](#)
- [Nutanix](#)
- [Pivot3](#)
- [StratoScale](#)

When evaluating these solutions, ask the vendor for a demonstration unit that can be installed in your office, so you can evaluate products alongside your existing environment. It is also important to remember that though HCI is based on concepts you may already know, the design and implementation of HCI depends on the vendor. Fortunately, most of the HCI vendors offer training to help you before and after you have made your selection.

Conclusion

This paper is an introduction to the world of Hyper-Converged Infrastructures and a new world of possible solutions to the ever-complicated art of data center management. While HCI may not be a panacea, it does an exceptional job of unifying virtual data center management, enabling data centers of all sizes to scale in affordable increments, and helping us take control over our applications' performance experience.

Learn More

To learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge, check out the following Global Knowledge courses:

[Cisco HyperFlex: Implementation and Administration](#)

[Nutanix Enterprise Cloud Platform Administration 5.0](#)

[SDN Fundamentals: The Future of Networking](#)

[SDN Planning Workshop: Demystifying Vendor Solutions](#)

[VMware vSphere: Install, Configure, Manage \[V6.5\]](#)

[VMware NSX: Install, Configure, Manage \[V6.2\]](#)

[VMware Horizon 7: Install, Configure, Manage \[V7\]](#)

[VMware Virtual SAN: Deploy and Manage \[V6.2\]](#)

Visit www.globalknowledge.com or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

Brian Eiler, VCI (L2), VCIX, PCNSE, ASE, CCNA is an instructor at Global Knowledge where he teaches Cloud, VMware, and SDN courses. Brian's technical background includes nearly 20 years of implementing, teaching, and writing about products from vendors including VMware, Palo Alto Networks, Cisco, Microsoft, Citrix, IBM, HP, Dell, and EMC. In addition to his other certifications, he holds VMware's advanced instructor credential (VCI-Level 2) and the VMware NSX advanced implementation credential: VCIX. He has taken his classroom and field experience to the masses by authoring a number of publications, training courses, and instructional videos. Brian is based in Fort Wayne, Indiana.