



HYPER-CONVERGED INFRASTRUCTURE: REFLECTION OF MY 10 YEARS JOURNEY



What is Hyper-Converged Infrastructure

Hyper-Converged Infrastructure (HCI) is a type of IT architecture that integrates compute, storage, networking, and virtualisation resources into a single, unified system. This system is typically managed through a centralised software platform, making it simpler and more efficient to deploy, manage, and scale IT resources compared to traditional, siloed data centre architectures.



Key Components of HCI

Compute: The processing power, typically provided by servers, needed to run applications and workloads. Today's HCI offers a wide range of processor architecture from x86 to ARM based designs. In an HCI environment, compute resources are tightly integrated with storage and networking components, all within the same physical nodes or appliances. This integration allows for a more efficient, scalable, and easy-to-manage infrastructure.

Storage: The space where data is stored, often provided through a combination of solid-state drives (SSDs) and hard disk drives (HDDs) within the same appliance. Unlike traditional storage systems, where storage hardware (like SANs, NAS, or DAS) comes bundled with its own proprietary software, HCI uses Software-Defined Storage (SDS) architecture which decouples these layers, allowing the storage management software to run on commodity hardware. This approach provides greater flexibility, scalability, and cost efficiency in managing storage resources.

Networking: The connectivity that allows communication between the different components within the HCI system, as well as with external networks. Most HCI would be using Software-Defined Networking (SDN), which is a network architecture approach that separates the network control plane (which determines where traffic is sent) from the data plane (which forwards traffic to the selected destination). This decoupling allows for centralised control and more flexible management of network resources, enabling dynamic, programmable network configuration rather than relying on traditional, hardware-based methods.

Virtualisation: A software layer that abstracts and pools the underlying hardware resources, allowing multiple virtual machines (VMs) to run on the same physical hardware. Virtualisation is a foundational technology in HCI, playing a critical role in how HCI systems consolidate and manage compute, storage, and networking resources.



How HCI Works

In an HCI environment, the traditional three-tier architecture (separate servers, storage, and networking equipment) is collapsed into a single, integrated appliance. This is made possible by software that tightly integrates these resources, enabling them to be managed as a cohesive unit. The software-defined nature of HCI allows for automation, better resource utilisation, and easier scaling, as new nodes can be added to the system to increase capacity and performance.

Benefits of HCI

There are many reasons to consider HCI over traditional bare-metal build up for a modern data centre or server room. Some of the benefits of adopting HCI are:

Simplicity: HCI reduces complexity by consolidating multiple components into a single system that can be managed through a unified interface.

Scalability: HCI systems are highly scalable; organisations can start small and expand their infrastructure as needed by adding more nodes.

Cost Efficiency: By eliminating the need for separate storage and networking equipment, HCI can reduce capital expenditures and operational costs. This is most evident in the cost of traditional storage infrastructure.

Improved Performance: The integration of compute and storage resources within the same system can lead to lower latency and improved performance for applications.

Ease of Management: With centralised management, IT teams can monitor and manage the entire infrastructure more efficiently, reducing the time and effort required to maintain it.



History of Hyper-Converged Infrastructure

Early Pioneers in HCI

The term “Hyper-Converged Infrastructure” was coined by Gartner in 2012. Gartner introduced the term to describe an emerging technology trend where compute, storage, and networking resources are tightly integrated into a single, software-defined system. But even before the term was introduced, there were a few groups of people working on the concept. In the late 2000s, a couple of start-ups stood out.

1. SimpliVity: SimpliVity, founded in 2009, was one of the early pioneers in the HCI space. The company set out with a mission to simplify IT infrastructure by integrating compute, storage, and networking into a single, scalable appliance. SimpliVity's OmniCube became known for its ability to provide advanced data efficiency, including deduplication and compression, across the entire IT stack. This innovation was crucial in reducing data centre complexity and costs, helping organisations streamline their IT operations. SimpliVity was eventually acquired by Hewlett Packard Enterprise (HPE) in 2017, further integrating its technology into HPE's ecosystem.



2. VCE (Vblock): VCE, a joint venture formed in 2009 by Cisco, EMC, and VMware, introduced the concept of converged infrastructure with its Vblock systems. Unlike a typical HCI design, which combines compute, storage, and networking into a single appliance, VCE's Vblock offered a pre-integrated, validated combination of servers, storage, and networking equipment. VCE's approach targeted large enterprises looking for rapid deployment and simplified management of IT infrastructure, setting the stage for what would eventually evolve into hyper-convergence. Though not an HCI solution by strict definition, VCE's impact on the IT landscape laid the groundwork for future developments in infrastructure simplification. VCE could more accurately





be termed as “Engineered Systems” not unlike what Oracle did with their Exadata platform.

3. Nutanix: Founded in 2009, Nutanix is often regarded as the company that truly brought HCI to the forefront. Nutanix introduced the concept of a software-defined, hyper-converged infrastructure with its NX-series appliances. What set Nutanix apart was its ability to decouple the underlying hardware from the management software, allowing enterprises to build highly scalable and flexible data centres. Nutanix’s Acropolis Hypervisor (AHV) further disrupted the market by offering an alternative to VMware’s dominance, enabling organisations to manage their entire infrastructure stack through a single, unified platform. Nutanix remains a key player in the HCI market, continuing to innovate with its hybrid cloud solutions.



New Players in the HCI Landscape

1. Huawei FusionCube: Huawei entered the HCI market with its FusionCube solution, targeting a broad range of enterprise customers. Huawei’s approach focused on providing a highly integrated, all-in-one infrastructure solution that combined compute, storage, networking, and cloud management capabilities. FusionCube’s integration with Huawei’s broader cloud ecosystem allowed it to offer a seamless hybrid cloud experience, catering to enterprises looking to bridge their on-premises infrastructure with public cloud environments. With its strong presence in the Asia-Pacific region and beyond, Huawei has become a significant competitor in the HCI space, challenging established players like Nutanix and VMware. Huawei also offers their own Kunpeng processor architecture (ARM based) as an alternative compute solution.

2. Other Emerging Players: In addition to Huawei, other players like Cisco HyperFlex, Dell EMC VxRail, and VMware vSAN have also made significant strides in the HCI market. These companies have leveraged their existing strengths in networking, storage, and virtualisation, respectively, to offer robust HCI solutions that cater to various enterprise needs.



My Journey with HCI

My journey with HCI began in 2015, during a critical period when I was in an organisation was looking to replace the aging data centre servers. As part of this initiative, we explored the three major HCI vendors at that time (SimpliVity, VCE and Nutanix) to determine the best fit for our needs.

The process involved an extensive proof of concept (POC) phase, where we rigorously tested the capabilities and limitations of each vendor's offering. During this evaluation, we discovered that SimpliVity, while innovative, required special PCI card connectivity. This requirement posed a significant limitation, as it restricted our hardware options and added complexity to the deployment.

On the other hand, VCE offered a robust and integrated solution but was found to be too rigid in its design. The lack of flexibility made it difficult to adapt the infrastructure for various workloads, which was a critical consideration for our organisation's evolving needs.

Ultimately, we decided to go with Nutanix, which provided the right balance of flexibility, scalability, and ease of management. Nutanix's ability to seamlessly migrate most of our VMware workloads onto its platform allowed for a smoother transition and better alignment with your goals. This decision marked the beginning of a significant transformation in your data centre operations, leveraging the full potential of HCI to drive efficiency and agility in our IT infrastructure.

After choosing Nutanix for our HCI platform, we initially continued running VMware on Nutanix. This was a natural choice since our IT operations team was well-versed in VMware, ensuring a smoother transition and minimising disruption. However, as Nutanix's Acropolis Hypervisor (AHV) matured and demonstrated increased stability and robustness, we saw an opportunity to optimize further.

Recognizing the potential for significant cost savings, we made the strategic decision to migrate all our VMware workloads to AHV. This move not only simplified our infrastructure by reducing the reliance on multiple platforms but also resulted in substantial licensing savings.



This transition allowed our organisation to leverage the full capabilities of Nutanix while maintaining the performance and reliability needed for our critical workloads.

One of the significant challenges I encountered during my journey with HCI was the complexity and cost associated with Oracle licensing on virtualised hardware. Oracle's licensing model in a virtualised environment often presents challenges that can lead to increased costs and operational complexities. This issue became a major concern as it directly impacted the efficiency and cost-effectiveness of our infrastructure.

To address this challenge, I made a strategic decision to migrate all our Oracle workloads to PostgreSQL. This migration allowed us to move away from the restrictive and expensive Oracle licensing, enabling a more flexible and cost-effective solution. The transition to PostgreSQL not only resolved the licensing issues but also aligned with our broader goals of optimising your IT infrastructure and reducing costs. This move further solidified our commitment to leveraging open-source technologies and modernising our data centre operations.

As I move on to other organisations in my journey, I remain a strong believer in HCI and continued to use Nutanix as the platform of choice. We experimented with various hardware vendor and licensing models too. I started my Nutanix journey with Dell, but soon found the limited range of compatible server nodes to be a challenge in shifting our workload profiles. In subsequent organisation, I have also use Nutanix on Lenovo and even trying out their TruScale model. Lenovo TruScale is a flexible, subscription-based licensing model that provides infrastructure as a service (IaaS). It allows organisations to access and utilise Lenovo's hardware, software, and support services on a pay-as-you-go basis, like how cloud services are billed. While it brought us closer to how public cloud providers bill us, it did not align with our internal workload profile and thus while it was a novel idea, it was not a good fit for us.

In my latest deployment, I decided to go with Nutanix's own NX series of servers (OEM from SuperMicro). As x86 infrastructure is increasingly commoditised, the choice of server made little difference in terms of performance. Additionally, the NX series comes pre-installed and



fully tested by Nutanix at factory, saving us a lot of work in reimaging and testing on deployment.

As we move to a hybrid-cloud architecture for robustness of our infrastructure design, we have a goal of making our private cloud on HCI as close to the public cloud as possible. One of the key add-ons we deployed is Nutanix Flow. **Nutanix Flow** is a software-defined networking (SDN) solution integrated into the Nutanix HCI platform. It provides advanced network security, micro-segmentation, and visibility features that help us secure our workloads and improve network management. This key feature allowed us to implement security groups like we do in the public cloud. This was a big game changer for us as we no longer had to design the network security differently for cloud or on-prem solutions.

As we also move to the use of more containers and Kubernetes workloads, we are also exploring **Nutanix Karbon**, which is a Kubernetes management platform that enables us to deploy, manage, and scale containerised applications on Nutanix HCI. While not perfect, it gives us an option to remove virtual machines and deploy containers into our HCI.

The Future of HCI

As HCI continues to evolve, the next logical step appears to be the integration of HCI with hybrid and multi-cloud environments. The rise of edge computing, artificial intelligence, and machine learning workloads is also driving the need for more advanced, distributed HCI solutions that can deliver low-latency, high-performance capabilities at the edge of the network.

Moreover, the concept of disaggregated HCI is gaining traction, where the traditional HCI stack is broken down into its constituent components, allowing enterprises to scale compute and storage independently. This disaggregation provides greater flexibility and efficiency, especially in large-scale, heterogeneous environments. This is very clear in Huawei's HCI strategy where they are pushing FusionCube (Hypervisor) and FusionCube (Storage) as a disaggregated design.



Composable Infrastructure is another trend on the horizon, where infrastructure resources are treated as a service, dynamically allocated to meet the needs of various applications. This approach aims to deliver cloud-like agility and simplicity within on-premises data centres, offering a glimpse into the future of enterprise IT infrastructure.

Conclusion

In conclusion, while HCI has made significant strides in simplifying and transforming data centre operations, its evolution is far from over. The integration of HCI with emerging technologies and the shift towards more dynamic, service-oriented infrastructure models will likely define the next chapter in its history. I look forward to the future of more innovation infrastructure designs that would allow us to be a lot more agile to meet the demands of the business and application evolutions.