Modernizing Data Centers: The Era of Hyper Convergence

Abstract

The Business 4.0™ era requires organizations to operate intelligent and agile systems integrated in the cloud. These systems generate an astronomical amount of data of extensively diverse range. This drives the need for flexible infrastructure that can be scaled quickly to meet the ever-increasing demand for data storage. What is also needed is the processing power to run complex machine learning algorithms to generate meaningful insights, cost effectively. Achieving it is possible only if data centers actively evolve with the changing industry needs to support both modern and legacy workloads without compromising on performance and security.

This paper will discuss how hyper converged infrastructure (HCI) contributes to solving the data center problem and the essential aspects that merit consideration when selecting an HCI solution.

The Data Center Challenge

Cyber systems require a data center infrastructure that can aid cloud computing, Internet of Things (IoT), big data, and other emerging technologies. To harness the full potential of digital innovations, business leaders are coming up with digital transformation strategies and IT divisions are opting for software-defined resources and scale-out platforms. This is where hyper converged infrastructure is gaining popularity. Like other infrastructure models, HCI can also be sourced from various original equipment manufacturers (OEM) in different configurations and needs to be introduced carefully in the data centers.

Selecting the Best Fit

Before looking into the factors that should be considered when selecting an HCI solution, here are a few use cases for understanding where and how it can deliver maximum benefits:

i. **Virtual desktop infrastructure (VDI):** Every VDI instance is processed and stored by the server, and even a relatively small number of instances can demand significant computing resources and network access. As the company grows, more HCI nodes can be added to the existing cluster. Deploying HCI for VDI hence provides easy scalability that siloed infrastructure lacks.

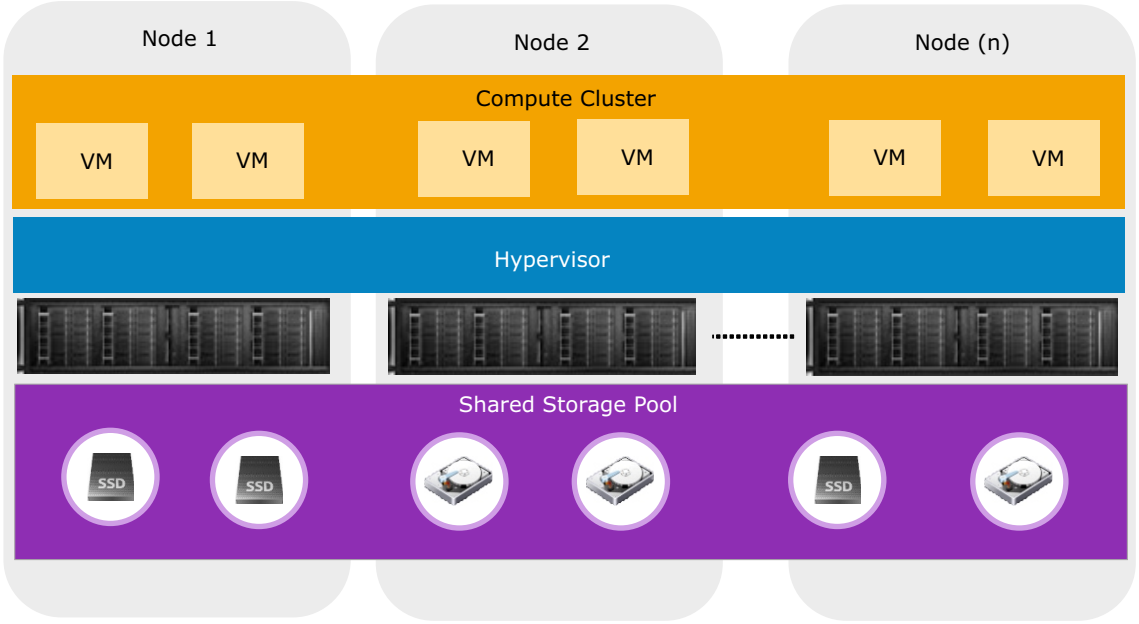

Figure 1: Logical view of HCI cluster

ii. **Test and Dev:** It is common knowledge that test and development environments should be similar to the production environment for test runs to depict the actual behavior of an application. Such environments therefore must be highly dynamic, as they may require an increase in capacity over time. In conventional environments, adding storage to the network is a challenging task, whereas the modular design of HCI allows the addition of nodes to increase the capacity as and when the need arises.

iii. **Multi-access edge computing (MEC)[1]:** In MEC, telecom service providers deploy HCI at cellular base stations to offload traffic from the network backbone. This requires computing at the edge (such as next to a highway) and micro data centers (which have less than 20 racks) across locations. Such an HCI should have faster CPUs, GPUs, RAMs, network, and storage input/output, but can run on low storage capacity as a lot of data is discarded after analytics and some of it is saved in the cloud as well.

iv. **Enterprise applications:** HCIs are being used to host enterprise applications such as RDBMS, CMS, and mail servers. These applications are very high on storage requirement. The subsequent addition of storage and computing resources increases the complexity of the infrastructure, making the applications underperform. For seamless scalability and streamlined performance, systems that are predictable and easier to deploy are preferred. In this regard, HCI checks most of the boxes in the expectation list. It is easy to configure and scale, and performs consistently when supported by flash storage.

v. **Big data analytics:** Extracting meaning from data quickly is critical in today's dynamic business landscape, and that necessitates rapid adaptation of the infrastructure for big data. Distributed file systems, NoSQL databases, and data analytics applications are the most popular big data workloads. Such workloads have very specific infrastructure requirements. For instance, Hadoop replicates data across nodes, which requires twice or thrice the storage capacity. Similarly, NoSQL requires rapid data processing due to the varieties of formats it deals with while data analytics applications require faster access to data to extract insights quickly. Compared to traditional infrastructure, HCI is easier to manage and scale even in clustered environments.

Factors to be considered when selecting an HCI solution:

- **Deployment architecture:** Broadly, there are two types of HCI available in the market – software-centric (build your own server or BYOS) and hardware integrated (turnkey). BYOS provides flexibility with the hardware components; turnkeys are quick to configure and expand the data center. Therefore, customers who want to continue with their preferred hardware vendor may opt for a BYOS-oriented HCI and incorporate the certified appliances from their hardware supplier into it. Alternatively, they can buy turnkey solutions from another HCI vendor, where they do not get the freedom to choose hardware vendors. In both scenarios, product support is provided by the HCI vendor.

- **Virtualization:** A variety of hypervisors is supported for virtualization. VMware vSphere is the popular choice while MS Hyper-V and Citrix Xenserver are also great options. There are a few solutions that offer license-free hypervisors leveraging open-source Kernel-based Virtual Machine (KVM), which helps reduce the cost while providing similar functionality.

- **Hardware configuration:**

i. **CPU** – HCI solutions run virtualized workloads, and to speed-up instruction execution in a virtualized environment, CPUs should have integrated virtualization extension. To support increased virtual memory (VM), it is crucial to have more computing power, which can be obtained by increasing the number of sockets or cores. Even though a dual-core CPU may not be as fast as two single-core CPUs, it is advisable to select multi-core CPUs to save cost, which is generally based on the number of sockets, though it may vary for different OEMs.

ii. **Memory** – Memory capacity starting from 64 GB to 3 TB can be used as long as it is proportional to the number of VMs helping in better resource utilization and the number of slots available for scaling the memory.

iii. **GPU** – When the workload requires running machine learning algorithms, simulations, or analytics, it is advisable to opt for high-end GPUs that assist in math-heavy software.

iv. **Storage** – HCI solutions can be hybrid (hard disk and flash) or all-flash. For business-critical applications, where minimal latency and high throughput is required, all-flash storage is advisable. For other cases, hybrid flash can be used. SSD and NVMe SSD are available for flash storage; the latter is the

latest innovation and preferred for workloads that require higher bandwidth and queue depths, such as databases.

- **Data availability** – HCI solutions have built-in erasure coding (EC), which protects data at the cluster level. In disaster-like situations, it is imperative to have a comprehensive policy-based data protection strategy, including the creation of multiple backup copies replicated across geographies.
- **Efficiency** – Techniques like de-duplication and compression offer the flexibility to have disk capacity smaller than the size of data. However, though effective, it may impact the system performance during execution since there is no dedicated CPU for running such programs. Such services are generally not provided as default but as additional licenses.
- **Manageability** – Resource administration, granular monitoring of critical system parameters (such as capacity and performance), and centralized management console is provided with most of the HCI solutions. Besides, analytical insights to predict future resource usage and integration with public cloud platforms to operate in Hybrid cloud environment are also essential.

Powering Data Center Evolution with HCI

With the proliferation of digital technologies, it is also critical that data centers, which serve as the backbone of the digital revolution, rapidly evolve with the change. Data centers are, therefore, undergoing a paradigm shift – from a hardware-centric approach to an application-centric one. This has enabled organizations to fully manage their data center engagements from a centralized console. When evaluated with the right set of technologies for workloads, operational expenditure on HCI is lower than in case of legacy infrastructure, with lesser data center footprints and reduced management overhead. Since the onset of its innovation, HCI has seen major enhancements from leading vendors and has now become a mainstream technology. Gartner predicts that HCI is expected to grow at 48% CAGR by 2021[2]. This further highlights its proven capability to take the business forward when deployed judiciously.

References

1. SDxCentral, The Advantages of HCI for MEC, accessed February 21, 2019, https://www.sdxcentral.com/cloud/converged-datacenter/definitions/hci-mec

2. Gartner, Driving Opportunity with Hyperconverged Integrated Systems, accessed February 21, 2019, https://www.gartner.com/webinar/3753375

About The Authors

Pawan Ratwani

Pawan Ratwani is a solutions architect with TCS' HiTech business unit. He has over 14 years of experience and has played several key roles in engineering and high tech domains. Ratwani specializes in storage manageability and data protection and is responsible for digital solutions leveraging data storage technologies. He has a Bachelor's degree in Information Technology Engineering from Dr. B.R. Ambedkar University, UP, India.

Sanjeev Kumar

Sanjeev Kumar is a domain consultant with the Software Product Engineering (SPE) group of TCS' HiTech business unit. He has over 14 years of experience and leads pre-sales and technical solution development in the areas of servers, storage, and manageability. Kumar has extensive experience in embedded system technologies and storage device protocols such as firmware design, NAND flash-based devices development, NVMe protocol, and 'system on chip' (SoC) implementation. He has a Bachelor's degree in Computer Science Engineering from the Madan Mohan Malaviya Engineering College, UP, India.

Contact

Visit the Hitech page on www.tcs.com

Email: hitech.marketing@tcs.com

Subscribe to TCS White Papers

TCS.com RSS: http://www.tcs.com/rss_feeds/Pages/feed.aspx?f=w
Feedburner: http://feeds2.feedburner.com/tcswhitepapers

About Tata Consultancy Services Ltd (TCS)

Tata Consultancy Services is an IT services, consulting and business solutions organization that delivers real results to global business, ensuring a level of certainty no other firm can match. TCS offers a consulting-led, integrated portfolio of IT and IT-enabled, infrastructure, engineering and assurance services. This is delivered through its unique Global Network Delivery Model™, recognized as the benchmark of excellence in software development. A part of the Tata Group, India's largest industrial conglomerate, TCS has a global footprint and is listed on the National Stock Exchange and Bombay Stock Exchange in India.

For more information, visit us at www.tcs.com

Experience certainty. IT Services
Business Solutions
Consulting

TCS Design Services I M I 03I 19