Hyperconverged Infrastructure Evaluation Guide

December 2020

Evaluator Group

Enabling you to make the best technology decisions

Evaluation Guide Overview

The Evaluator Group Hyperconverged Infrastructure Evaluation Guide is part of a series of Evaluation Guides designed to help evaluate storage technology alternatives. These documents are for IT professionals seeking a neutral, objective discussion of the design considerations behind new products, technologies and trends. Evaluator Group Evaluation Guide series are not vendor-sponsored, but offered as part of a paid subscription. They are developed based on our review of technology options and strategic analyses of how they can best be used in today's information storage environments.

With this information Evaluator Group recommends IT professionals review their requirements, starting with this guide and the associated materials available, and use our team of experts to help make selections. This Evaluation Guide plus the Hyperconverged Comparison Matrix, Product Briefs and Product Analyses, are designed to assist potential buyers in understanding the options and products available and to help match requirements to the available technology choices.

The Hyperconverged Infrastructure (HCI) Evaluation Guide will briefly explain what a hyperconverged infrastructure is and then detail its advantages and disadvantages in a typical IT environment to suggest when to consider deploying this technology. Then we will discuss specific features that help differentiate the many products available in this space and provide some useful evaluation questions.

What is Hyperconverged Technology?

A traditional, physical datacenter infrastructure is a connected ecosystem of servers, networking and shared storage systems. These storage systems are accessed by multiple servers over a physical network specifically designed to carry data using block (FC or iSCSI) or file (NFS, SMB, etc.) protocols. Server resources are shared among clients (users) by a physical LAN or WAN. This shared resource flexibility enables administrators to adapt more easily to changing business requirements than with legacy direct-attached storage.

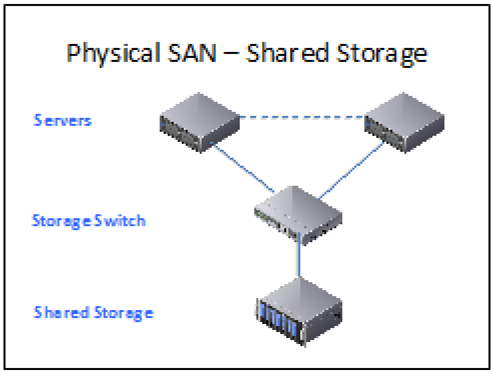

Figure 1: Physical Data Center Infrastructure

The requirement by IT organizations to provide both higher levels of service and improved storage efficiency are driving a re-architecting of IT environments. In this ongoing push for improved infrastructure agility, IT organizations are embracing the concept of "convergence", the combining of servers, storage and often networking functions in an effort to simplify infrastructure design and speed deployment. The first products were called "Converged Infrastructures" (CI), essentially bundles of traditional IT components sold by the rack. Hyperconverged Infrastructures (HCI) take this concept a step further by physically combining these functional components into the same server chassis, providing a modular architecture that starts smaller and scales more easily than CI deployments.

Hardware Architecture

Hyperconverged Infrastructure is a term used to describe what some vendors are calling "infrastructure in a box". Since storage devices and storage networking connectivity are designed into a single server chassis coupled with pre-installed virtualization and management software, this all-in-one solution is often referred to as an "appliance". HCIs are also including Graphics Processing Units (GPUs) as options for higher performance and analytics-heavy use cases. (see Figure 2).

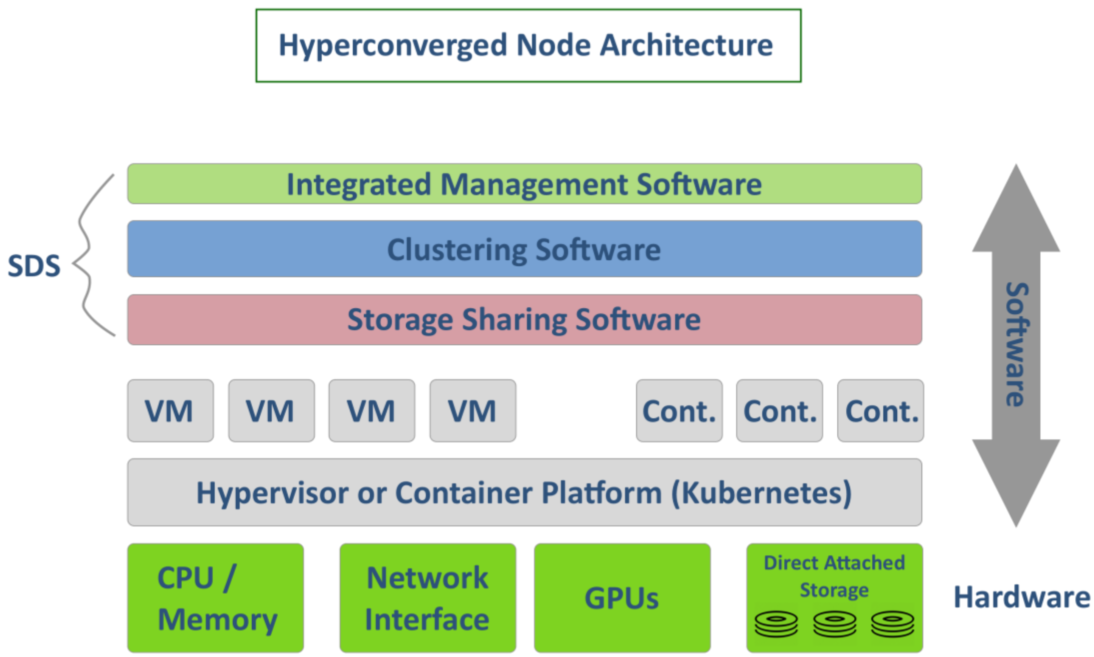

Figure 2: Hyperconverged Architecture

Originally, HCI nodes contained relatively fixed amounts of storage and compute (with memory) capacity, which meant that scaling the cluster would add both resources simultaneously. As clusters have grown and HCI deployments become more specialized, vendors have come out with more node configurations, some storage-heavy and some compute-heavy models. But most HCI architectures still require that each node run a hypervisor (for the software stack – see “Software Architecture”), which means they generate a certain amount of compute overhead, even for a node designed to provide storage capacity, and they consume a hypervisor license. Similarly, when only compute expansion is needed, most HCI nodes still have at least a minimum of storage capacity.

Disaggregated Architecture

Disaggregated architectures are a little different, enabling the use of truly dedicated storage or compute nodes in the cluster and helping to address this resource inflexibility to control costs. In a disaggregated architecture compute nodes don’t have any storage capacity (outside of a boot drive), but do have larger CPUs, more memory and often GPUs (see Figure 3). On the other hand, storage nodes have minimal CPU power and memory capacity, enough to run their virtualization, clustering and management functions, but no GPUs and no hypervisor (see Figure 4). Some disaggregated architectures, like Microsoft Azure Stack HCI (with Storage Spaces Direct SDS) or Dell EMC VxFlex, can also be configured with *both* storage and compute functions on the same nodes, in a “combined” architecture, the way most HCIs have traditionally been designed. Disaggregated HCIs handle their SDS software stack differently than traditional HCIs as well (see “Software Architecture”).

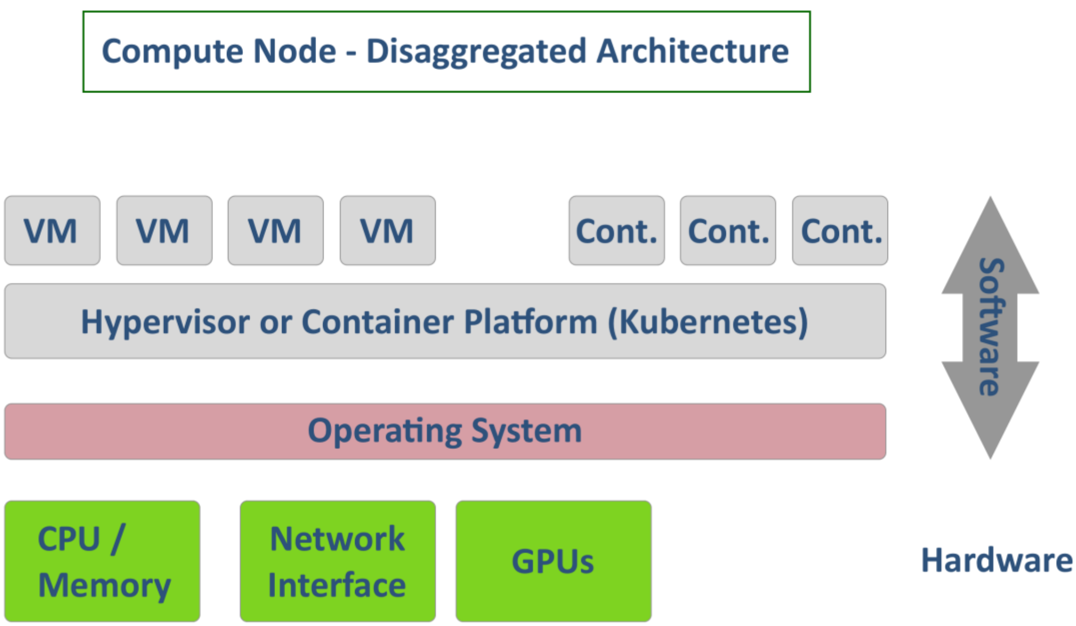

Figure 3: Compute Node Architecture

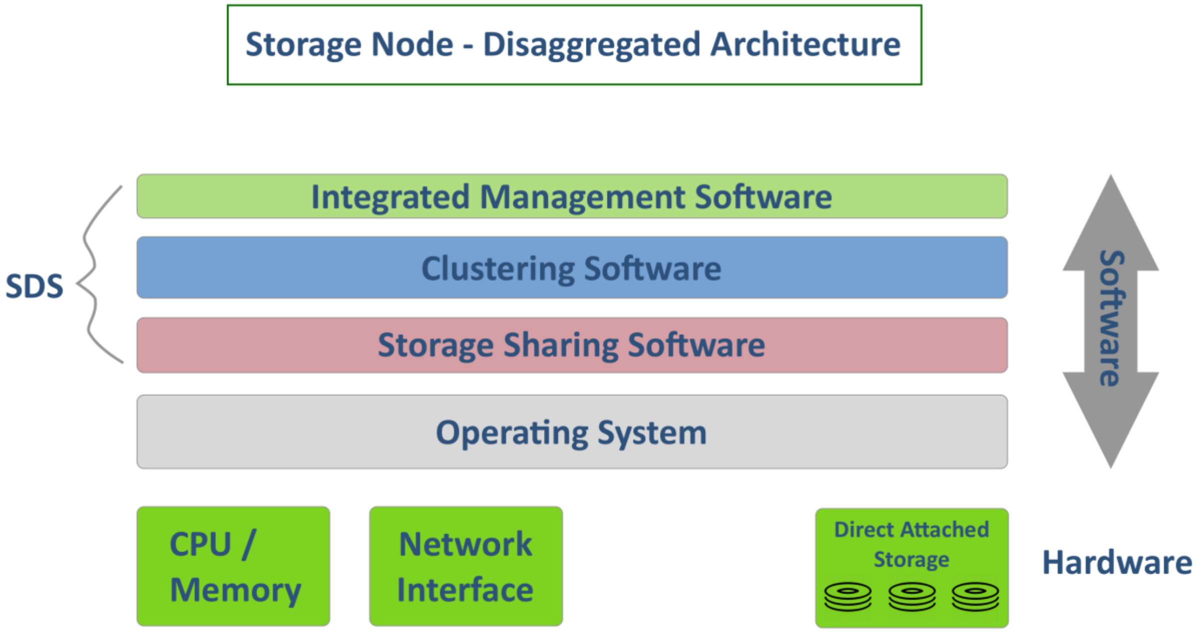

Figure 4: Storage Node Architecture

Software Architecture

HCIs have pre-installed software-defined storage (SDS) that creates a virtual storage layer – block- or file-based - between the storage devices internal to each node, or occasionally direct-attached. (See Figure 2). This enables storage to be shared between nodes, resulting in better capacity utilization and more flexible scaling. By design, virtualization manages where data is located and controls the access to that data for users and applications. SDS also manages the creation and placement of redundant data across the cluster as required for the appropriate level of availability and data protection.

Where traditional HCIs run this SDS layer as a VM on each node, disaggregated architectures typically run the SDS software natively on the OS, or in the OS kernel as Windows Storage Spaces Direct does with Azure Stack HCI. Similarly, VMware vSAN runs in the VMware kernel. In disaggregated architectures, Hypervisors and/or Container platform software runs on compute nodes, to support workloads, but not on storage nodes (see Figure 3.) For storage nodes, SDS runs on the operating system as well, but doesn't run on compute nodes at all (see Figure 4).

Evaluator Group Comment:

Running the SDS software natively on a storage node OS, or embedding it in the hypervisor or the operating system kernel reduces software layers that can impact storage performance,

compared with running SDS as a VM. There are also be cost advantages to eliminating the hypervisor as disaggregated architectures can do.

Networking

Storage networking has historically been ignored by the HCI vendors (except Cisco HyperFlex originally). Adding switches and networking complicated the simple deployment that these products emphasized, and, many HCI use cases involved relatively small clusters. However, the maturing HCI market is increasing the demand for more comprehensive "rack level" solutions, with higher capacity clusters and dedicated networking infrastructure. More HCIs are now including switches and networking management software, as well as integration with software-defined networking solutions such as VMware NSX or Cisco AHI.

Deployment

Current drive capacities give HCIs the ability to scale very large, into the hundreds of TBs per node for some vendors' products, creating the potential for clusters of a petabyte or more, depending on the storage clustering technology. However, in practice, most HCI deployments are more limited in scope. In the SMB and some mid-market companies, HCI clusters are providing a comprehensive data center infrastructure, replacing traditional SAN and server components, typically in the tens of TBs. In larger companies and enterprises, HCIs were originally used to support specific projects like VDI, test and development or to provide a standard IT infrastructure for remote offices. As the technology has matured enterprises are using HCI clusters to support more core data center use cases as well, especially infrastructure consolidation.

Evaluator Group Comment:
In research and regular contact with IT organizations, the Evaluator Group is finding that larger organizations are acknowledging the value of the the HCI approach, with most expecting to eventually have HCI deployments in place, including "mainstream" IT environments. Many want the flexibility to choose the hardware platform (such as one they have already standardized on) and are opting for HCIsoftware solutions that are offered through multiple server vendors.

Why Hyperconvergence?

Virtualized infrastructures have become common in most IT organizations. But there is a level of complexity that goes with building a virtualized environment that could make rapid deployment difficult. HCI automates that deployment, centralizes management and operations and can scale out performance and resources incrementally, up to the maximum size of the particular SDS cluster.

Virtualization has proven valuable in helping to increase utilization of resources and allow their dynamic allocation, with web-scale capabilities in some cases. Most hyperconverged products offer the ability to distribute a virtualized infrastructure over many remote offices and disaster recovery sites, while providing centralized management.

Hyperconverged Infrastructures provide a number of benefits, the most significant being simplification and agility. Compared to traditional IT infrastructures – storage networks, shared storage arrays and compute servers - HCIs are easier to set up, easier to run and easier to expand, reducing the time required to deploy new application infrastructures and the overhead to keep them running. This means HCIs can even be deployed and managed by the departments consuming these resources, such as server virtualization teams, not just by dedicated IT staff. This simplicity also makes HCIs an easy solution for dedicated infrastructure to support a specific IT function, such as disaster recovery, data protection or remote offices.

While HCIs speed infrastructure design and deployment, they do so at a cost. The software required to virtualize storage and provide resiliency in a cluster can add significantly to the cost of the servers and storage devices that comprise each HCI node. For enterprise IT professionals HCIs can mean less control, less sophisticated management and lower levels of security, compared with the traditional server/SAN environments they are familiar with. In some environments this is a good trade off, but in others it's too inflexible. It is important that the potential disadvantages to an HCI decision be considered as well as the advantages.

The following is a list of potential advantages and potential disadvantages of deploying an HCI cluster. It should be noted that not all of these apply to every product or to every deployment.

HCI Potential Advantages

- Improves IT agility - set-up time is reduced and configuration flexibility is increased
- Reduced IT infrastructure complexity – simplifies design decisions and reduces operational overhead
- Makes virtualization easier to implement – "VM-centric" solution is ideal for virtual server or VDI infrastructure
- Improved resource efficiency – can provide better utilization of compute and internal storage than traditional IT infrastructure, especially with a disaggregated architecture
- Centralized management of compute and storage – can provide a cost-effective, stand-alone compute solution for small offices and branch locations
- Simple off-site solution – HCI clusters can be easily set up in remote locations for backup and DR
- Enables user-based operation – simplicity can allow non-IT groups to manage their own infrastructures in some environments, reducing IT overhead
- Provides standardization of IT infrastructure – common compute, storage and storage networking functionality can be used for many applications and use cases

- Enables more efficient scalability – incremental, node-based expansion can be more efficient than traditional server/SAN infrastructures, and disaggregated architectures can be even better
- Provides effective cloud infrastructure - as a flexible, scalable, comprehensive compute solution HCI appliances can be used to support private and hybrid clouds

HCI Potential Disadvantages

- Vendor lock-in - turnkey appliance format forces users to buy additional modules from the same vendor
- Lack of resource independence – requirement to add compute and storage together in HCI nodes can lead to sub-optimal resource utilization – although disaggregated HCI solutions can help
- Technology change rate - differences between server, storage and storage networking may cause additional work in updating or cause older, less capable technology to be retained long after useful life
- Doesn't easily use existing storage assets – use of existing shared storage systems may be difficult or sub-optimal
- Management and control – HCIs can lack the sophisticated tools and management software that enterprise IT is accustomed to
- Security – consolidating management and control responsibilities can reduce security oversight and less sophisticated management tools can reduce visibility
- Less budget flexibility – Combining storage and compute (and often networking) resources forces companies to make these purchases at the same time, instead of spreading them out.

Evaluator Group Comment:
This list is meant to provide general guidance about the potential benefits and shortcommings of an HCI solution. Actual experieinces will vary, based on worklaod, HCI configuration, even the expertise of the IT personnel involved.

HCI Features

In order to determine what type of system is best suited to meet specific business objectives and technology requirements, the technical aspects of the HCI solution should be considered. Most vendors offer similar levels of compute and memory resources. The real product differentiation is in platform choices, storage architecture choices (disaggregated, combined or both), storage hardware, storage software features/functionality, hypervisor support and overall system management. These characteristics are included in the Hyperconverged Comparison Matrix and described below as they are listed in the Matrix.

Hardware OEM

Hyperconverged Infrastructure vendors use industry-standard x86 server hardware and "commodity" storage components to keep costs down and simplify the development process, leaving most of the innovation to the software. That said, many HCI vendors make their software available for sale with other vendors' hardware. This could be a factor to consider if a particular hardware OEM is important. Or, a software-only "roll your own" HCI solution may be an alternative (see Evaluator Group Research on Software-Defined Storage")

Models Available

While most HCI vendors use the same basic hardware architectures, different products offer different selections of models. CPU options can vary as well, but are typically less of a differentiator. Most vendors offer configurations with more CPU cores and memory for compute-centric workloads and storage-heavy configurations for capacity-centric workloads. Total storage capacity, hard drive and flash (hybrid) capacity and all-flash options can vary significantly between vendors, as can the availability of "storage-only" nodes. Overall, vendors chosen should offer enough model options to support current use cases and projected growth.

HCI Software

HCIs include a software-defined storage program that virtualizes physical storage on each node and creates a shared pool of capacity. This software typically runs as a virtual machine but can also be run natively on a server, and one, vSAN, is actually part of the VMware hypervisor. Some products are also integrated with an infrastructure or hypervisor management platform or have other management options.

Appliance Architecture

This section reveals whether an HCI has a traditional or disaggregated architecture – or can have nodes in either/both configurations. It also provides details about HCI clustering that can be important if rack space is limited or scaling capacity is important. These include the number of nodes in each server appliance (1, 2 or 4), the size (1U or 2U), the minimum and maximum number of nodes supported and how they support adding storage capacity. The OEMs that sell a vendor's HCI software is also listed.

Compute and Networking

Each node contains one or two CPU sockets with a range of cores and memory capacity available. Each physical chassis also offers a finite number of data ports and control ports. These specs are listed for each model.

Caching

HCIs, like traditional storage systems, employ various caching methods to maximize performance. These details are listed, as well as the type of storage device used for the cache and whether it's NAND flash or DRAM.

Storage Supported

This section lists number of individual storage devices (hard drives and SSDs) and total storage capacity each node can support. For hybrid models that capacity is comprised of hard drives, with SSDs used for the cache. In all-flash models capacity is comprised of SSDs. Most vendors support Graphics Processing Units (GPUs) on at least one model.

Data growth is a fact of life for IT, making the ability to scale the infrastructure essential, especially storage. In most hyperconverged systems storage and compute resources are tied together into modules for easy implementation. But this convergence of resources can limit options for efficient scaling of storage, since adding hyperconverged nodes to increase capacity also increases compute power.

As this category matures, vendors are starting to address this need for independent scaling. Some are simply creating storage-heavy nodes that increase the ratio of drive capacity to compute power. Other companies enable external storage arrays or virtual SAN storage nodes to be connected into the hyperconverged cluster. And a few HCIs support separate storage and compute nodes in a disaggregated architecture.

Scalability is a complex metric, one that involves more than just raw storage capacity. In the virtualized environments where HCI appliances are deployed, the number of virtual machines or virtual desktops supported can be an especially pertinent measure of scale. For this reason some vendors use tools such as **IOmark** to establish accurate per-VM or per-VDI metrics.

Hypervisor Support

Many companies use more than one hypervisor and most HCI appliances support more than one hypervisor as well, so vendor choice should take the hypervisor into consideration. Some products offer a proprietary hypervisor option which can save money by eliminating licensing fees and provide some additional VM mobility features. Microsoft includes Hyper-V hypervisor along with Storage Spaces Direct in the Windows Server operating system. Support for containers is also becoming more common.

Storage Management

This is one of the biggest strengths and potential values of hyperconverged solutions. Management functionality can make life easier for the IT administrator but also benefit their internal customers by enabling faster response to new requests and changes. Most products offer a dedicated management interface plus an API plug-in for VMware or Hyper-V that integrates these hypervisors' management interfaces into the HCI cluster. All HCI solutions include monitoring and management functions, but some offer more advanced analytics capabilities that provide trending and operational guidelines for resource allocation and expansion.

HCI products often provide self-service portals and interfaces for local and remote users to meet their infrastructure needs based on the policies and controls set by the administrator. Ideally, the centralized management software automates most tasks and policies for ease of use. Multi-tenancy and Quality of Service are management features that HCI appliances are starting to include to better support these kinds of environments.

Some HCIs also organize management around VMs. This means users or admins specify things such as data resiliency, backup frequency, retention, the performance they need, etc., for each VM and the system handles the details. Whether implemented automatically through policies or set manually, VM-level management can enable non-IT personnel to successfully manage an HCI, or just save a lot of effort for IT admins.

Data Resiliency / Protection

Data resiliency is the method used to assure that data can survive system failures – typically hardware failures like a drive loss – without losing data. All HCI solutions have a fundamental architecture for protecting against data loss from failures at the node level (failure of an appliance) and/or at the component level (failure of a storage device). Some products use a “shared nothing” architecture and distribute data services and metadata across all nodes while others simply create copies of data objects and store them on alternate nodes.

Most hyperconverged solutions offer RAID options (simple mirroring, single or double parity) or a forward-error correction scheme (“erasure coding”). These methods provide different levels of resiliency and different efficiencies in terms of storage capacity and data handling overhead.

Data protection refers to the creation of data copies and storing them locally or remotely, a capability that most HCIs have. Some include additional data protection options that automate the entire process and can replace the need for traditional backup or disaster recovery. Most are integrated with one or more third party backup applications as well.

Replication

Most HCIs provide the ability to take snapshots (metadata-based point in time copies) and clones, and replicate those copies to other nodes in the cluster, as well as asynchronously replicate those copies to a remote cluster. Many support synchronous replication (stretched clusters) providing high availability data protection as well. This ability to get copies off site provides the foundation for a disaster recovery process either by restoring data from the remote copy or actually running VMs from the DR site or from the cloud (see Cloud-based Features).

Data Reduction / Capacity Optimization

Deduplication, compression, thin provisioning and data coalescing technologies have come a long way in the past decade and most HCI appliances offer several of these features. While they could be thought of as data services, processes like data reduction[1] can change the economics of an HCI appliance decision by increasing the "effective" capacity of the system significantly. Care should be taken to understand how data optimization processes will impact effective capacity for intended workloads and how those processes will affect performance

Cloud-based Features

Most companies have adopted a hybrid cloud model in which an on-premises HCI cluster connects with one or more public clouds in order to provide different IT functions, such as backup, DR, cloud-native development, etc. This section includes five features, native to the HCI, that involve connecting with the public cloud, moving data between the on premises cluster and one or more public clouds, managing data and applications in the hybrid cloud and performing advanced analytics on operational data sets.

Cloud Backup - This service uses the public cloud as a backup target, with software running in the cloud to act as a remote replication point for the on-site HCI cluster, using the native replication capability that is included in the SDS layer. Restores involve replicating data back from the cloud backup repository. Since most backups are done at the VM level, cloud backup is sometimes combined with disaster recovery (see below). Capacity for cloud backup is sold by the HCI vendor or by one of the public cloud vendors (most often AWS, Google or MS Azure)

Cloud-based DR - Similar to cloud backup, cloud-based DR requires capacity in the public cloud to store critical data and applications (typically VM images), software functionality to restart them (failover) in the cloud, and often the ability to "fail back" when the primary site can be brought back on line. DR uses the ability of maintain synchronized copies of these data on-premises and in the cloud. Cloud-based DR is frequently sold as a service (DRaaS) by the HCI vendor.

Cloud Native - This service involves setting up and running application development in the public cloud in a container environment (typically Kubernetes). It also includes the ability to move applications between the cloud and the on-premises HCI cluster and often between public clouds.

Cloud Orchestration - This feature provides simplified or automated "self-service" deployment, management or operations of private and or hybrid cloud infrastructure. This includes private cloud (on-premises) hardware and software, and VMs or containers in a hybrid cloud environment, by IT or end-users.

---

[1] Evaluator Group Data Reduction Modeling tool estimates impact of deduplication and compression http://www.evaluatorgroup.com/data-reduction-estimator/

Cloud-based Analytics - This feature involves predictive analytics (often using AI/ML techniques) of a cloud-based data set, captured through installed base telemetry and shared with the HCI user. It is used to manage and optimize resources, monitor and alert for critical conditions and to create insights for business intelligence.

Other Advanced Features

Auto-tiering – the capability in hybrid systems to store more frequently used data on flash to improve application performance

Data Encryption – data-at-rest protection that's provided by the drives themselves (self-encrypting drives, or SEDs) or implemented as a software-based service.

Quality of Service – policy-based controls that are designed to ensure performance for specific VMs

Multi-tenancy – the ability to isolate workloads on the same physical storage volumes

WAN optimization – techniques that minimize the bandwidth required to remotely replicate a given data set

Evaluation Questions

This Evaluation Guide has focused on the technical and functional aspects of hyperconverged technology and of the available HCI appliance solutions. The questions below are designed to help decide whether an HCI appliance is a better solution than a traditional IT architecture and how to identify which HCI products warrant more detailed evaluation.

- For my environment, what are the advantages of an HCI solution compared with a Converged Infrastructure, or a traditional server-SAN/NAS, and do they outweigh the disadvantages?
- How will my company save money by deploying a hyperconverged datacenter infrastructure? What efficiencies will the company gain and how will they be measured?
- Will there be any business impact (positive or negative) to implementing this technology in the datacenter? What personnel changes will it induce?
- Do the HCI solutions being considered support the server hardware my company has standardized on?
- Will the HCI cluster be set up and/or run by IT generalists or even by non-IT personnel?
- Do the HCIs being considered provide the appropriate management tools?
- How many virtual machines or virtual desktops will this infrastructure need to support, and do the products being considered provide per-VM or per-VD metrics needed to ascertain their scalability?
- What levels of data reduction should be expected, based on intended workloads, and how does that affect capacity requirements?

- What are the requirements to scale, based on the defined application requirements (net of data reduction)?
- What levels of resiliency will be needed and what costs in terms of storage capacity and overhead will this generate?
- Will this system need to integrate with existing infrastructure, in the data center environment or in the cloud? If so what are the requirements? If not, what impact will that have?
- Will the data protection features of the solutions being considered meet my application business needs or what 3rd party applications do they support?
- What data services (features) are important?
- What cloud-based services are needed and which public clouds are or will be used?
- What will be the process for handling technology updates that occur on different schedules for servers, networks and storage?
- Is it possible to disable the hyperconverged infrastructure if problems arise?

Evaluator Group EvaluScale™

Working with many IT clients, Evaluator Group has developed a list of the most important criteria for making product selections. These criteria and the associated requirements comprise the EvaluScale. For each product, Evaluator Group publishes a Product Brief that includes an EvaluScale showing how the product measures up. Requirements do vary depending on usage and IT environments, variations that generally follow a segmentation of high-end enterprise, mid-range or entry-level. The EvaluScale incorporates these differences into each requirement and orders the criteria based on Evaluator Group's opinion and information gathered from IT client engagements.

	Criteria	Description	Requirement	Features / Functionality Examples
1	**Performance**	Resources and design elements to maximize performance	Standard perf metrics yet to be established. Offers adequate flash capacity and some perf accelerators - NVMe, disaggregated arch a plus	Adequate flash capacity for hybrid and AF configs, may include automated tiering, DRAM caching, high-core CPUs, data locality features, GPU options
2	**Economics**	Effective cost of system and operation, incl data reduction	Have characteristics designed to reduce TCO, incl data reduction, data resiliency, admin overhead, licensing, HW costs, including disaggregated arch	Dedupe, thin provisioning, compression, erasure coding, parity RAID, simple design, hypervisor and container options, adv/automated features policy/VM-based management, 2-node or single-node ops, impact of OEM/supplier relationships on cost

3	**Scalability**	Total storage capacity, flexible, efficient expansion	Provide node and cluster capacity for expected growth, scale needed resources easily, disaggregated resource options a plus	Variety of device sizes, types, density (drives/node), seamless, flexible expansion, external storage connectivity options
4	**Management Options**	Hypervisor mgmt integration, internal/ISV mgmt pkg, automation, VM-level mgmt	Have mgmt integration with hypervisors, internal mgmt sw or 3rd party mgmt pkg - adv analytics capability, self-service/automation a plus	API or plug in for vCenter, MS Sys Center, Cisco UCS, dedicated internal storage mgmt or integrated w/ ISV mgmt pkg, resource analytics or "SRM" functions, DC automation, EU self-service ops, policy-based mgmt by VM
5	**Advanced Features**	Leverage functionality enabled by SDS, provide unique or adv efficiency or mgmt features	Have some features beyond basic data services - can incl QoS, external data sharing options, adv data and mgmt efficiency, automated processes, etc	Dedicated QoS and multitenacy features, external storage sharing (block/file), encryption, WAN optimization, application awareness, automated storage services (BU, DR, etc), policy-based mgt, data locality, adv efficiency features
6	**Vendor/Product Stability**	Company viability, product longevity, impact of OEM partnerships	Established company with stable partnerships and mature product technology	Start up or established IT vendor, does product rely on OEM relationships that could change, does a substantial installed base exist, does company offer comprehensive services and support
7	**Data Protection / Business Continuity**	Device- and node-level failure protection, BU/DR options, stretched clusters, HA, failover	Have multiple resiliency (RAID) options, DP function, remote synch mirroring - automated BU, HA failover, efficient recovery a plus	Local and distributed RAID, mirroring and parity, forward error correction (erasure coding - EC), one-touch BU, stretched clusters, automated or policy-based BU/DR
8	**Platform Options**	Hypervisor support, container support, related features	Have support for multiple hypervisors and some integration with container platform- migration, orchestration, adv mgmt features a plus	Support KVM, Hyper-V, Xen, support container mgmt/orchestration layer, integrated storage features for containers, migration utility for hypervisors or containers
9	**Cloud Options**	Private cloud features/functionality, public cloud connectivity for hybrid deployments	Have comprehensive strategy for hybrid cloud including min of 3 out of following 5 cloud-based features: Backup, DR, Cloud Native/Kubernetes support, Orchestration and Cloud-based Analytics	Connectivity with Azure, AWS and/or Google clouds for BU, DR/DRaaS. Native Kubernetes support on prem with migration/management in cloud, Self-service features for IT and end users, advanced analytics leveraging installed base
10	**Model Options**	Variety of available appliance models and configurations	Vary compute-to-storage ratio with range of resource combo options, nodes per chassis, entry/ROBO model, support for HCI and disaggregated arch (dedicated storage and compute nodes) a plus	Models for: General purpose, ROBO, VDI, storage-centric and high-performance workloads, multiple AF configurations, single and multi-node chassis, single CPU option

Summary

Hyperconverged Infrastructure solutions have become very popular as companies strive to improve their IT organizations' abilities to deploy infrastructure quickly and make changes easily. The success of these appliance-based products has created a market with over a dozen viable vendors, now mostly major IT infrastructure suppliers, but also a few smaller companies. This guide and the associated Evaluator Series Research materials are designed to help companies confirm if an HCI is the right technology for their environment and then to make the best choose between the vendors' products that are available.

There are many possible benefits to deploying an HCI solution, including simplifying IT infrastructure, centralizing resource administration and lowering total costs. Hyperconverged Infrastructures have the potential to help IT organizations deliver on these benefits; but there are tradeoffs, such as a reduction in vendor choice and less configuration flexibility.

When comparing HCI vendors, companies should focus on technology attributes that differentiate the available solutions. In terms of hardware, these include model configuration options, scalability (how large and small) and OEM platforms available. If scaling efficiency is important, a disaggregated architecture should be considered. On the software side, data resiliency, optimization and synchronous replication (stretched clustering) are important, as are data protection, data management and cloud-based features.

About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. **www.evaluatorgroup.com** Follow us on Twitter @evaluator_group*