

Key Criteria for Evaluating Enterprise HCI

TABLE OF CONTENTS

- 1** Summary
- 2** Key Criteria Report Methodology
- 3** Considerations About the Evaluation Criteria
- 4** Table Stakes
- 5** Key Criteria
- 6** Critical Features: Impact Analysis
- 7** Near-Term Game-Changing Technology
- 8** Conclusion
- 9** About Enrico Signoretti
- 10** About GigaOm
- 11** Copyright

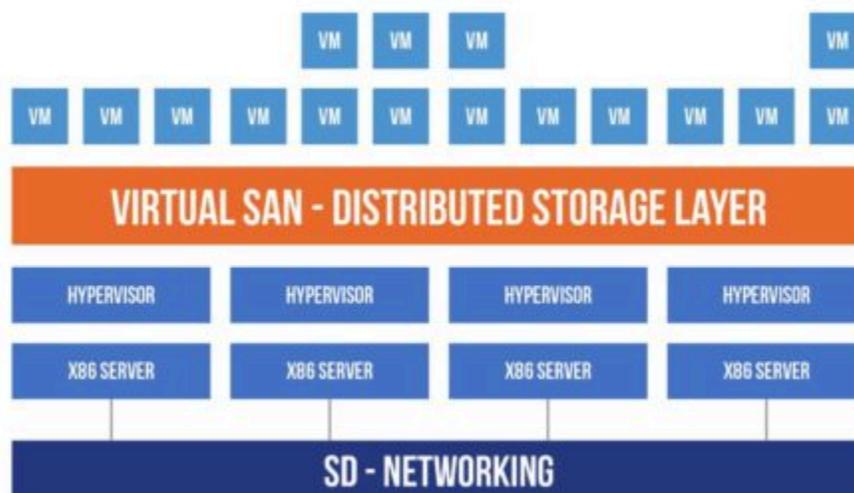
1. Summary

IT organizations of all sizes are facing several infrastructure challenges. On one hand, they repeatedly receive pressing requirements coming from the business to keep their company swift and proactive, while taking on new digital transformation initiatives. On the other hand, they struggle to keep their budget under control, provision new resources quickly, and manage the increasing complexity while maintaining efficiency at reasonable levels.

A cloud-only IT strategy is not a viable option for many organizations, now more are interested in hybrid scenarios from which they can get the best of both worlds. Unfortunately, by mixing cloud and traditional IT infrastructures there is a tangible risk of creating silos, going exactly towards the wrong direction, and complicating the overall infrastructure even more, introducing inefficiencies. Unlike traditional infrastructures, hyper-converged infrastructures (HCI) are a step forward in simplicity and flexibility. HCI enables organizations to hide IT infrastructure complexity while realizing the benefits of a cloud-like environment. HCI simplifies operations and makes the transition of data and applications between on-premises and cloud easier.

Usually based on commodity X86-based hardware and virtualization software, HCI is a software-defined solution that abstracts CPU, memory, networking, and storage devices, organizing them in resource pools. HCI allows the administrator to combine and provision these resources quickly in the form of virtual machines (VMs) and, more recently, as independent storage resources like network attached storage (NAS) filers or object stores (see Figure 1). Management operations are also simplified, allowing an increase in productivity levels of the infrastructure while reducing the number of operators and sysadmins per virtual machines (VMs) under management.

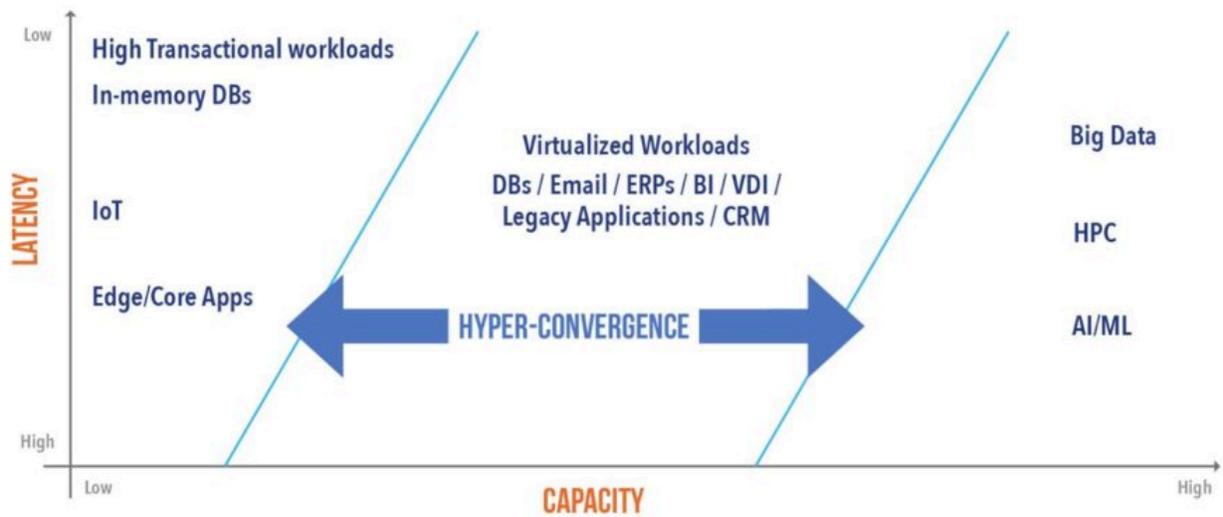
Figure 1. HCI Logical Scheme



Source: Gigaom (2019)

HCI, originally chosen by small- and medium-sized enterprises as an alternative to their traditional virtualized infrastructures for its simplicity and a very limited set of workloads, like VDI for example, is quickly expanding its reach to larger organizations while covering an increasing number of workloads, including all types of virtualized workloads, as well as big data analytics, edge computing use cases, and more. In fact, for the most advanced solutions, the trade-offs on efficiency imposed by simplicity are becoming less relevant and their feature set allows to address a very large set of applications now. See Figure 2 below for further reference.

Figure 2. Workloads covered by hyper-converged infrastructures from the data storage perspective



Source: Gigaom (2019)

HCI is a mature market now and solutions can be divided into three categories:

- **Enterprise solutions.** With a focus on large organizations' needs, they are characterized by an extensive feature set, high scalability, core-to-cloud integrations, and tools that go beyond traditional virtualization platform management and up in the application stack.
- **Small/Medium Enterprise solutions.** Similar to the previous category, but simpler and less expensive. The focus remains the simplification of the IT infrastructure for virtualized environments, with limited core-to-cloud integrations and solutions ecosystem.
- **Vertical solutions.** Designed for specific use cases or vertical markets, they are very competitive in edge-cloud or edge-core deployments, but usually, with a limited solutions ecosystem. Often not very scalable, but efficient from the resource consumption point of view, these solutions integrate open-source hypervisors, such as KVM for example, to offer end-to-end support and better prices.

A hyper-converged solution aimed at serving one of the categories described above should include:

- **Scalability:** HCI infrastructures are based on a scale-out design. Even though modern hypervisors don't have the limitations they once had, scalability limits can still be found on the distributed storage layer (or virtual shared storage) and in the tools aimed at managing the cluster and additional services. Scalable HCI systems allow the organization to manage fewer clusters and organize resources in larger pools, minimizing failure domains while improving infrastructure availability and overall system efficiency. A fundamental aspect to take into consideration when evaluating HCI infrastructures is the flexibility of the solution in terms of initial configuration, expansion, and additional data storage services. In fact, the ability to configure nodes of different sizes, HW generation, and purposes increases the system lifespan while enabling the end user to respond quickly to his organization's business needs with the right amount of resources and budget allocation.
- **Feature Set:** HCI features include the distributed storage layer and management tools with third-party integrations, common across the board. Then, depending on the main focus of the solution and reflecting the categorization defined above, additional features are available on practically every product in the market. Enterprise-focused solutions have a wider range of functionalities with the goal to provide an end-to-end infrastructure stack for enterprises. They expand their reach to data protection, secondary storage, analytics, cloud services, container support, and more. A large feature set is better for simplification of large data center infrastructure and overall total cost of ownership (TCO) but is usually more expensive. On the other hand, HCI solutions designed for small/medium enterprises (SME), cloud service providers, and edge-core/edge-cloud use cases are usually less expensive and are designed to provide the best combination in terms of efficiency and costs for the use cases where they excel.
- **Ease of use and usability:** One of the main intents of hyper-convergence is to reduce infrastructure friction and complexity. To do so, HCI products should provide tight integration between all the components involved (hypervisor, virtual machines, storage, network, and hardware), leveraging tools and user interfaces that are easy to use, and aimed at reducing the time necessary to perform sysadmin tasks and hiding the complexity behind the scenes. Commercial hypervisors usually offer all the necessary APIs for this type of integration and their management console can be extended with specific plug-ins and UI elements to manage other components of the stack. A different approach is to make the management console of the HCI solution more central, managing all the components from there, including the hypervisor. There are advantages and disadvantages for both approaches but especially with open source hypervisors or multi-hypervisor support, the latter is the most common. Another aspect of usability comes from the sophistication of the entire HCI solution for handling upgrades, system reconfigurations, node decommissioning or repurposing, and so on.
- **Flexibility:** Reducing infrastructure complexity shouldn't come at the expense of flexibility and efficiency. In the beginning, HCI was a solution for a limited number of workloads and applications but now, with the right solution, it is possible to cover almost all enterprise needs. HCI flexibility depends on architecture design and how it can take advantage of different resources at its disposal. The ability to manage nodes of different types and sizes, mix hardware of different generations and vendors, taking advantage of advanced load balancing mechanisms for IO operations, are at the base of it. Even more so, the ability to integrate external storage systems or

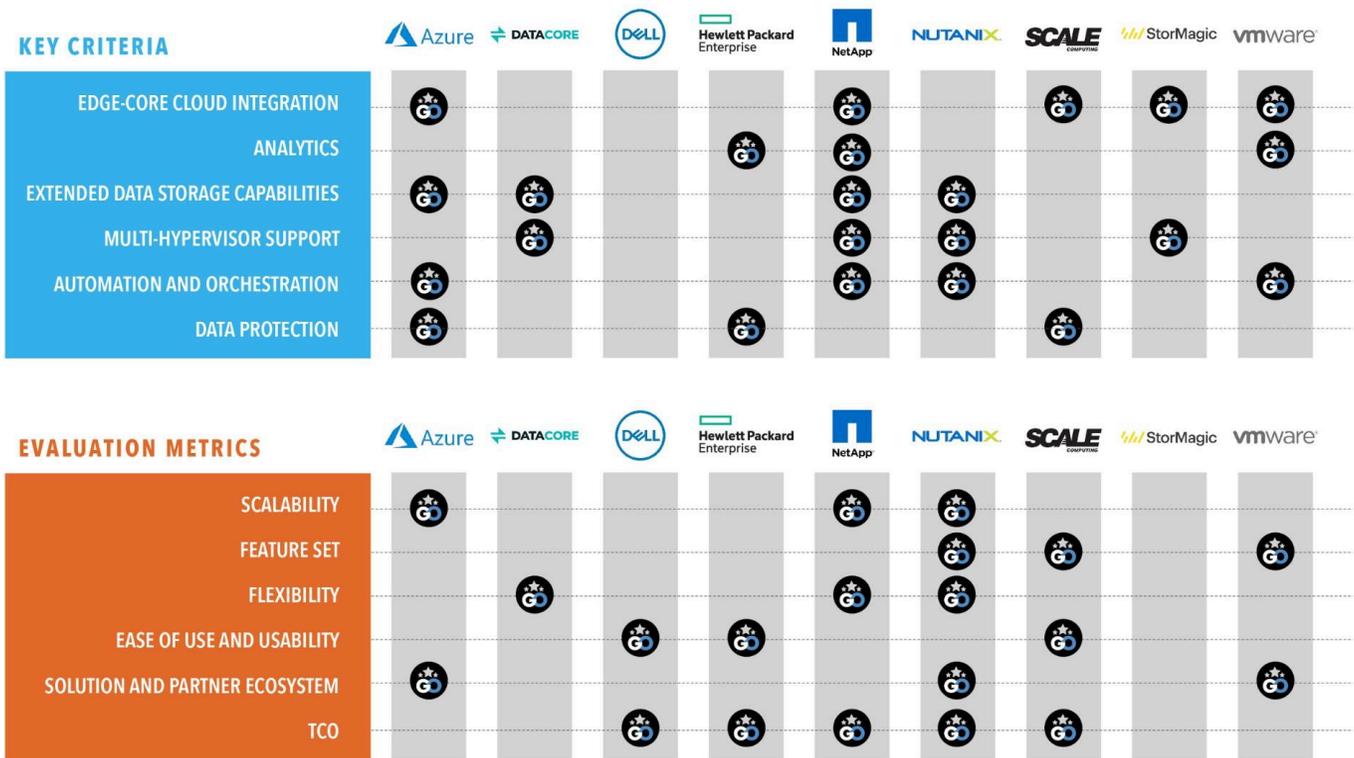
expose resources via block, file, or object interfaces to the rest of the infrastructure can be of great help to limit the initial investment, improve the return on investment (ROI) on HCI, simplify its adoption in heterogeneous environments, as well as create a smooth migration path from traditional to a hyper-converged infrastructure.

- **Solution and partner ecosystem:** The HCI stack is only a part of the entire IT infrastructure. It usually takes advantage of standard hypervisors, and integration with third parties can always be done at this level, but it is also true that the deeper the level of integrations, the better for the end-users. In fact, by leveraging plug-ins or APIs, a tighter integration can contribute to a better experience for the end-user. This is why many HCI vendors work with partners to build an ecosystem of solutions aiming at extending and improving the functionality of their HCI stack. Some vendors take advantage of the existing large product portfolio in areas like networking or storage, while others are building this ecosystem based on collaborations with other vendors. Areas of particular interest range from data protection to security and cloud.
- **Total Cost of Ownership:** Albeit HCI solutions are all designed to reduce complexity and therefore improve overall infrastructure TCO, some are more complete than others or with a better focus on specific use cases, allowing them to address TCO issues more efficiently. In general terms, TCO is very difficult to compute but when it comes to hyper-convergence solutions, the parameters described in the previous points all have an impact on TCO. Even more so, looking at the categories defined at the beginning of this chapter, the TCO of an HCI solution also depends on the architecture and its efficiency in running the workloads for which it has been designed.

Following the classification presented earlier in this chapter, the key metrics for the evaluation, and their impact we will analyze in this report, GigaOm has awarded Go Awards to most noteworthy solutions in the market. That said, in many cases, the evaluation of an HCI solution really depends on the specific user's needs and requires a thorough analysis of the use case. For this reason we organized the entire report to give information to the reader to help facilitate decision making.



GO Award Winners for HCI Capabilities, Innovation, and Key Metrics



2. Key Criteria Report Methodology

In a Key Criteria Report, we analyze several important features of a technology to better understand how they impact various parameters in relation to the needs of each IT organization. The goal is to group product features into three categories to explain the various characteristics and differentiators in modern technology. The goal is to help the end user understand the value of each capability presented by vendors, even those not mentioned in the document, helping to build mid to long term infrastructure strategy.

Technology solutions can be divided into three categories:

1. Enterprise solutions
2. Small/Medium Enterprise solutions
3. Vertical solutions

Even when a market has reached a very mature level, available technology solutions differ in their characteristics depending on several factors and how they can be integrated with existing infrastructures. Some vendors are more focused on medium to large enterprises while others are a better fit for smaller organizations or niche use cases. That said, in many cases, the evaluation of a technology is more dependent on the specific user's needs and not on the type of organization.

The most important metrics for the evaluation of a technology solution, aimed at serving one of the categories described above, should include:

- Scalability
- Feature set
- Ease of use and usability
- Flexibility
- Solution and partner ecosystem
- Total cost of ownership (TCO)

Considerations about the evaluation criteria: Selecting the right technology for an organization is not an easy task, and this is why we organized this document in three main sections:

- Table Stakes
- Key Criteria

- Near-term game-changing technology and features

Table Stakes

In this category, we pull together a series of system characteristics and features that are important when choosing the right solution. In most cases, these are just architectural choices that depend on the size of the organization, the features required, expected growth over time, and the type of workloads. Considering their maturity, the implementation of these features will not add any business advantage nor significantly change the TCO or ROI of the infrastructure.

Key Criteria

In this section, we discuss the features and the technology that really differentiate one system from another. In most cases, depending on real user needs, these functionalities can be beneficial and have a deep impact on one or more of the parameters mentioned. Therefore, details on the implementation are essential for understanding the benefits and how they apply to the infrastructure, processes, or business.

As soon as we move from what can be considered table stakes to critical and differentiating capabilities, aspects like architecture design and implementation regain importance and need to be analyzed in great detail. In some cases, the features described in the Key Criteria chapter are the solution core with the rest of the system is designed around them. This could be an important benefit for those organizations that see a real practical advantage in them, but it also poses some risks in the long term.

Over time, the differentiation introduced by a particular feature becomes less relevant and falls in the group of “table stakes,” while new system capabilities introduce new benefits or address new needs, with a positive impact on TCO, performance, flexibility and so on.

The system capabilities described in this section can bring several benefits to organizations of all sizes and with different business needs. This section is organized to give the reader a brief description of the specific functionality or technology, its benefits in general terms, and what to expect from a good implementation. In order to give a complete picture, we also include examples of the most interesting implementation currently available in the market.

Critical Features

We provide indications about the impact of features described in the previous sections on the metrics defined at the beginning of this report:

- The impact that each feature or capability has on each metric;

- The solutions in the market with the best impact on a specific metric at the moment of writing;
- The implementation which has the best impact on a specific feature or system capability.

The impact of the features on the selected metrics is not in absolute terms and should always be verified with the organization's requirements and use case. In fact, the weight we provide in the tables for the impact of the features on the metrics is for most common use cases in enterprise environments.

Near-term game-changing technology and features

In this section of the report, we analyze some of the most interesting products that are to be implemented in this technology in the next twelve to eighteen months. Some are already present in some form but usually as part of niche products or to solve very specific use cases. In either case, at this stage, implementations already available are not mature enough to be included in the group of the critical features. At the same time, when implemented correctly and efficiently, features based on this technology can really make a differentiation in the same metrics identified in this report.

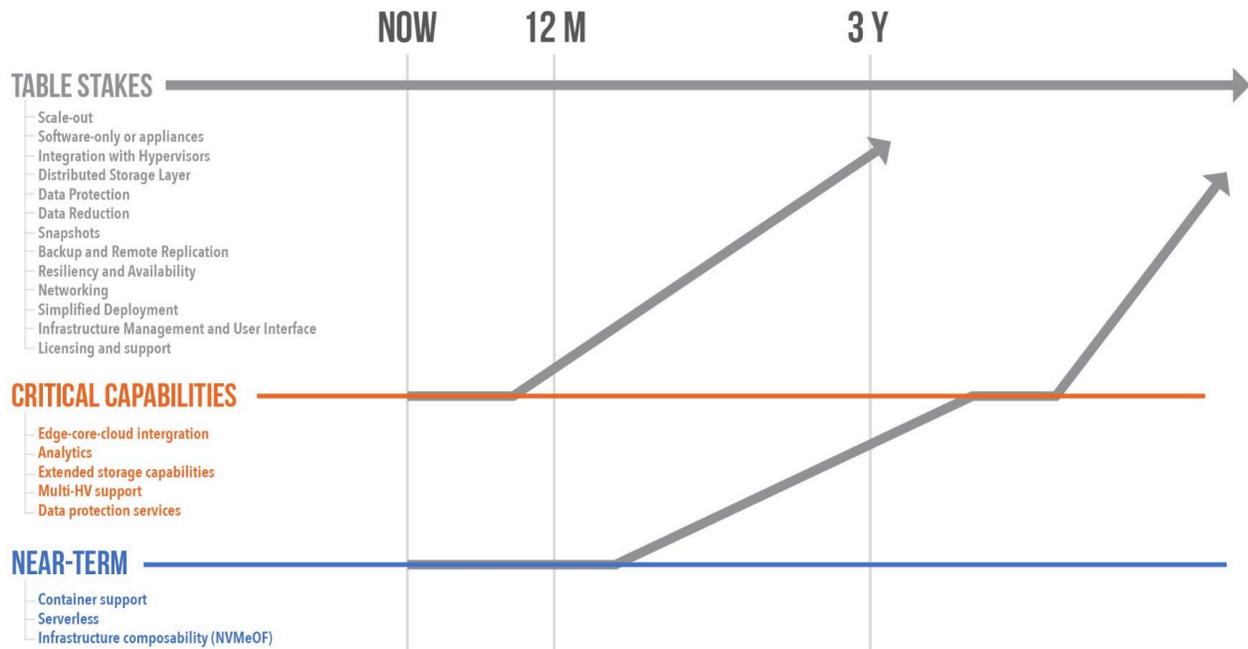
The goal of this type of report is to group features in three categories to give the reader the possibility of better understanding the various characteristics and differentiators.

3. Considerations About the Evaluation Criteria

Selecting a hyper-converged infrastructure for an organization is not an easy task, and this is why we organized this report in three main sections:

1. **Table Stakes:** The minimum feature set that is taken for granted in an enterprise-grade hyper-converged infrastructure system. Even though there are still differences among approaches and implementations of single features, most of the solutions have been maturing for years through multiple iterations and in most cases, they are architectural design choices that are well known and demonstrated where and when they best fit.
2. **Key Criteria:** In this section, we discuss the features and the technology that really differentiate one system from another. In most cases, depending on real user needs, these functionalities can be beneficial and have a deep impact on one or more of the parameters mentioned. Therefore, details on the implementation are essential to understanding the benefits and how they apply to the infrastructure, processes, or business.
3. **Near-term game-changing technology and features:** Most HCI systems have a lifespan range of about three to five years, and this is a market segment that evolves quickly. Looking at how technology, protocols, and system design are evolving will give organizations more information in order to find the best solution. Taking a look at the technology that is now in its early adoption phase, but expected to hit the market more aggressively in 12 to 18 months, will also enable the user to look at vendor roadmaps with a different and more critical eye.

Figure 3. Timeline: Impact of features and capabilities for the enterprise



Source: Gigaom (2019)

It is important to note that table stakes technologies and capabilities become less relevant and part of the feature set that everybody takes for granted in a relatively short time period. The same cycle continues for what it is now considered game-changing and innovative in HCI. Therefore, to get the best ROI from the hyper-converged infrastructure and increase its lifespan, it is important to always check what vendors can offer today and what they plan to release in the near future.

Systems that already have mature functionalities listed in the “Key Criteria” category can really make a difference in overall TCO, but it is also important to take into account all the specific needs of the end user, for which some of these features could be less relevant, hence less impactful. Much depends on the size of infrastructure and how critical the applications are that run on those systems.

Adopting new technology earlier is very beneficial, giving an edge over the competition, but this can also be a risk, depending on the maturity of the implementation. Even more so, some systems entirely designed around these innovative technologies could miss some of the features listed in other categories, reducing the number of use cases and compromising its usability for common workloads and scenarios.

Technology, functionality, and architecture designs that have demonstrated their value and sometimes made a fortune for a vendor are always adopted by others and they slowly lose their status of differentiator and competitive advantage. In the beginning, the design and implementation of these features were very important and at the base of technical discussions to understand the real value and

the trade-offs. But over time, with the maturation and evolution of these features, the differences become smaller, architecture designs tend to standardize, and they become part of the standard feature set of the HCI system.

An example to illustrate this process could come from basic data services such as VM cloning through snapshots at the storage level, a powerful tool to make copies of VMs quickly and efficiently. Introduced by a single vendor, in the beginning it was a huge differentiator, a feature that quickly found several use cases and making some applications much more efficient, for example, virtual desktop infrastructure (VDI). Other vendors rushed to build similar features and there was always a debate on limitations, constraints, or usability when the different implementations were compared. After several years, snapshots became just a checkbox in the features list. In fact, everybody agrees on what is really important in their implementation and how it should work. Some might be better than others, but they don't really make a difference except for very specific use cases.

4. Table Stakes

Scale-out

Scale-out (horizontal scaling) is a specific type of architecture design for which system scalability is obtained by adding additional units (nodes) with a similar amount of resources that are interconnected on a shared network. Almost all of the HCI solutions available in the market are based on this type of architecture, usually aligned with hypervisor requirements and best practices.

Modern hypervisors can scale to large cluster configurations and the limitations are usually found in the distributed storage layer. Typical configurations begin with a minimum of three nodes (two in particular configurations for edge use cases), with some solutions that can now manage dozens of nodes and capacities in the multi-petabyte range.

The complexity usually associated with scale-out infrastructures is abstracted by the hypervisor and its management console. Initial deployment of HCI-based solutions takes less time than traditional virtualized infrastructures.

Software-only or Appliances

Hyper-converged infrastructures are software-defined solutions that run on standard x86 servers. Some vendors leverage proprietary PCI cards to optimize deduplication operations but in many cases, this is an option and this card could be theoretically installed in any server with a PCI slot.

Organizations usually have two options:

1. **Reference Architecture, software-only:** The vendor provides the software and a compatibility matrix and the user can purchase hardware from the hardware vendor of choice. This method gives more freedom of choice and is usually preferred by very large organizations that already have framework agreements with hardware vendors, cloud providers, or for specific edge use cases where there are demanding hardware requirements in terms of power consumption, space, security, cost, and so on.
2. **Pre-configured appliances:** Many organizations, especially smaller ones, still prefer this purchasing modality. It simplifies the purchase process, installation, and support as well. Many HCI vendors partnered with several hardware vendors, which can now offer the end-to-end solution from their product catalog through certified resellers.

Licensing and subscription models reflect what is happening in the rest of the industry. Standard perpetual licenses and support plans are still available but practically all vendors are switching to subscription-based licensing models and in some cases, they also provide the hardware in a cloud-like fashion, on-demand. These new business models are aimed at helping customers transition their expenditure from capital expenditure (CAPEX) to operational expenditure (OPEX), limiting initial

investments and making infrastructure more flexible.

Integration with Hypervisors

To realize a virtualized infrastructure for the enterprise, there are a few options in terms of hypervisors: VMware vSphere, Microsoft Hyper-V, and one of the many variations of KVM (an open-source hypervisor for Linux).

VMware vSphere

VMware vSphere is the market leader and offers a broad range of functionality, most of which is exposed through APIs. At the same time, VMware also developed many other interfaces to gather information and interact with other components of the stack, enabling its partners to create seamless integration with their products and great user experience. This is also why the VMware partner ecosystem is the largest of the market and most of the HCI solutions are part of it.

Microsoft Hyper-V

Microsoft Hyper-V is part of Windows Server. It's a powerful hypervisor, and is now more focused on building hybrid architecture with Azure cloud in the backend than integrating with third-party tools. Therefore, the solutions ecosystem is centered around Microsoft Cloud solution partners to gain market share. Most of the solutions that support Hyper-V and Microsoft Azure Stack HCI are in the enterprise category, but Hyper-V also supports small configurations that can be deployed in edge use case scenarios.

KVM

Open-source KVM is gaining in popularity, but there are several versions of it, some of them directly maintained and supported by HCI vendors. Keeping track of different development branches and commercial distributions can be difficult, and this severely limits third-party support for this hypervisor. That said, a larger number of HCI vendors support KVM directly, using it in their solution as a cost-effective option or as the primary choice if they want to offer end-to-end solutions to their customers.

Many HCI solutions support more than one hypervisor, with the exception of a few vendors that choose to focus on a single virtualization platform because it targets a specific market segment of vertical use cases. In Figure 4 we represent how most common hypervisors compare on different applications.

Figure 4. Hyper-converged hypervisor use cases

	VMWARE VSPHERE	MICROSOFT HYPER-V	KVM
ENTERPRISE	★ ★ ★	★ ★	★
SMB	★ ★	★ ★	★ ★
EDGE / VERTICAL	★	★	★ ★ ★

Source: Gigaom (2019)

Distributed Storage Layer

The distributed storage layer provides the primary data storage service for the VMs and is one of the key components of every HCI solution. Usually presented as a virtual network-attached storage (NAS) or storage area network (SAN), depending on the exposed protocol, they have all of the data protection, availability, and resiliency features expected from these types of storage systems. Storage resources, flash-memory, and hard drives are usually installed on the same nodes that run the hypervisor, and are shared across the cluster.

Designed for niche use cases in the area of edge computing, most of these storage systems are scale-out, but not always as scalable as the hypervisor, limiting the size of the single cluster. There are three approaches to the distributed storage layer for HCI:

1. **Virtual storage appliance (VSA):** A VM that is managed by the same hypervisor that runs the rest of the VMs in the node. A VSA is more flexible, usually able to support multiple hypervisors, but this approach could result in higher latency.
2. **Integrated within the hypervisor or the OS:** The storage layer is an extension of the hypervisor, and it does not need some of the components (VM and guest OS) present in the previous approach. The tight integration improves overall performance, reports better telemetry about the workloads, and takes full advantage of hypervisor characteristics, but the storage layer is not portable.
3. **Specialized storage nodes:** The distributed storage layer is based on dedicated nodes with the goal of getting the best performance consistency and scalability for internal and external storage consumption. This approach, usually more expensive than the others in smaller configurations,

enables compute and storage resources to scale independently and granularly.

Performance

Usually, data storage performance is not one of the key design metrics when comparing HCI solutions, taking for granted that performance will be good enough for their needs and that they will always be able to add additional nodes if necessary. That said, the performance of HCI storage layer is usually good enough for all enterprise workloads, especially for systems listed in the third category. They can be flash-based or hybrid (flash + HDDs) and some of them can now also take advantage of storage-class memory to get better latencies.

Several mechanisms are also implemented by vendors to speed up IO operations; for example, the use of data placement techniques to keep data as close as possible to the VM that is accessing it. NVMe devices are becoming more common and supported by a growing number of vendors, again making IO bottlenecks a thing of the past.

Even though some vendors make use of hardware acceleration too, in the form of proprietary PCIe cards that are able to speed up some specific operations, it is pretty common that compute resources necessary for storage operations (CPU and RAM) come from the same nodes that run VMs.

Data Protection

All vendors now offer advanced data protection against multiple failures, including full node, single, and multiple-component issues. Distributed erasure coding protects data, balancing between performance and data footprint efficiency. This equilibrium is made possible by advanced instruction sets available in modern CPUs, new hardware like NVMe and storage-class memory (SCM) devices, as well as optimizations of the data path.

Some HCI solutions, especially those optimized for low-power CPU edge computing use cases, offer simpler protection mechanisms like n-way data replication, trading some storage efficiency for compute power and allowing users to build smaller infrastructures with fewer CPU power needs.

Data Reduction

Data footprint optimization is an important aspect of hyper-converged infrastructures. Deduplication, compression, and other techniques, like thin provisioning, can dramatically improve capacity utilization in virtualized environments, especially in use cases like VDI. Even more so, in order to optimize rack space utilization and get balanced servers, the number of storage devices installed in a single HCI node is limited.

On the flip side, data reduction is a CPU-intensive task that does not bode well with performance stability requirements of production applications. Vendors have different approaches to solve this challenge:

- **Post-processing:** After data is saved safely in a staging area or directly on disk, a process starts in the background and does the job with a low priority. This technique has a very low impact on latency visible at the front-end but is less efficient in space-saving and with heavy loads, this task could be stopped in favor of others.
- **Allocated CPU cores:** Some resources are always allocated to this task; the data optimization is performed inline without impacting applications but limiting the resources available to them.
- **Dedicated hardware:** A PCI card installed in the server with specific chips or dedicated storage nodes can perform these tasks inline without impacting any other task. Dedicated hardware provides predictable performance, but this also means that the servers have a PCI slot, which tends to make them larger and less power and space-efficient.

All of these optimization techniques have advantages and disadvantages. Post-processing is the most flexible, connecting to the cloud easily because there is no need for specialized hardware, but it is not the fastest. Dedicated hardware promises better performance, but it is difficult to apply in edge use cases and cloud installations.

Snapshots

All HCI storage systems have the ability to take VM snapshots or assist the hypervisor in taking them. The most sophisticated solutions enable creation and orchestration of VM groups to take application-consistent snapshots by quiescing data on persistent storage prior to taking the snapshot. Many administrators use snapshots as a first form of backup or they integrate with backup vendors to speed up backup operations. HCI snapshots can be useful for specific workloads, like virtual desktops, or to orchestrate system upgrades, software distributions, and other similar tasks.

Backup and Remote Replication

In addition to managing snapshots, HCI capabilities are increasingly becoming capable of managing remote replication for disaster recovery. Snapshots can be moved to a second HCI cluster geographically or sometimes to the cloud. Retention policies can be applied to make them expire automatically, keeping capacity utilization under control. Remote replication is now pretty common, allowing replication of a single VM or a group of VMs remotely. Depending on the implementation, it can be done by taking snapshots and syncing them remotely, taking advantage of features available in the hypervisor or using specific mechanisms implemented in the storage layer. The most enterprise-focused solutions offer additional tools for disaster recovery (DR) orchestration and data/copy management aimed at simplifying testing and DR processes for large environments.

With the increasing interest for edge solutions, backup and remote replication is an aspect of HCI that many enterprise organizations look at with a renewed interest. Many backup vendors have tools that can work seamlessly with virtualized and hyper-converged infrastructures, giving the organization the freedom of choice about the solution to implement, especially when the goal is to keep cross-platform compatibility and multi-hypervisor/cloud support.

Resiliency and Availability

HCI is inherently resilient because the hypervisor has many high availability mechanisms in place. HCI systems offer redundant components and features that improve overall system availability, including non-disruptive system upgrades, hardware replacement, and so on. Call-home function is standard on most HCI solutions particularly when the system is provided as an appliance. Even more, standard monitoring features are now joined by advanced analytics tools that the end users can use to spot anomalies and potential issues.

Alongside traditional remote replication, most advanced enterprise solutions can be configured in stretched clusters across multiple data centers, increasing overall infrastructure availability while simplifying operations and getting business continuity in case of a disaster.

Networking

Hypervisors have several components to build fully virtualized infrastructures, including virtual network interface cards (NICs) and switches. VMware goes even further with its software-defined networking (SDN) product, NSX. Usually, HCI solutions will not add more to this except for a few tools to simplify the management or configuration of external components like switches. This means that most of the networking capabilities found in HCI systems are based on the hypervisor's functionalities and are identical for all the vendors.

Simplified Deployment

All HCI solutions provide simple and automated deployment processes that allow the system to get up and running in a short time. This includes initial system configuration, storage layer, networking, and hypervisor setup. In many cases, the same tools are used also for cluster expansion and node decommission.

Infrastructure Management and User Interface

HCI products have their own management console, but many of them integrated well with the hypervisors that are part of the stack. Depending on the solution and the type of hypervisor, there are at least two levels of integration between the hypervisor console and the HCI solution:

1. **HCI Console manages the hypervisor/s:** This solution is pretty common among HCI products that can manage multiple hypervisors or KVM-based hyper-converged infrastructures in enterprise environments. This helps give the end-user a standardized view of the entire infrastructure and simplify operations. In the case of KVM, this approach is also a necessity because there is not a standard cluster management console, as there is for VMware and Microsoft products.
2. **HCI doesn't manage the hypervisor/s:** This is common among simpler solutions that have the goal to be more cost-effective and more focused on SMB or edge computing use cases. Most of these

solutions provide plug-ins and can be managed from the hypervisor's UI, allowing the hypervisor's management tools to take control of the entire HCI stack and manage it from within a single interface, simplifying day-to-day management tasks.

In many cases, HCI vendors offer a combination of these two approaches, allowing the administrator to choose what is best to manage daily operations.

Licensing and Support

The HCI market is following a trend that is pretty common in the entire IT industry now, with standard perpetual licensing replaced by subscriptions.

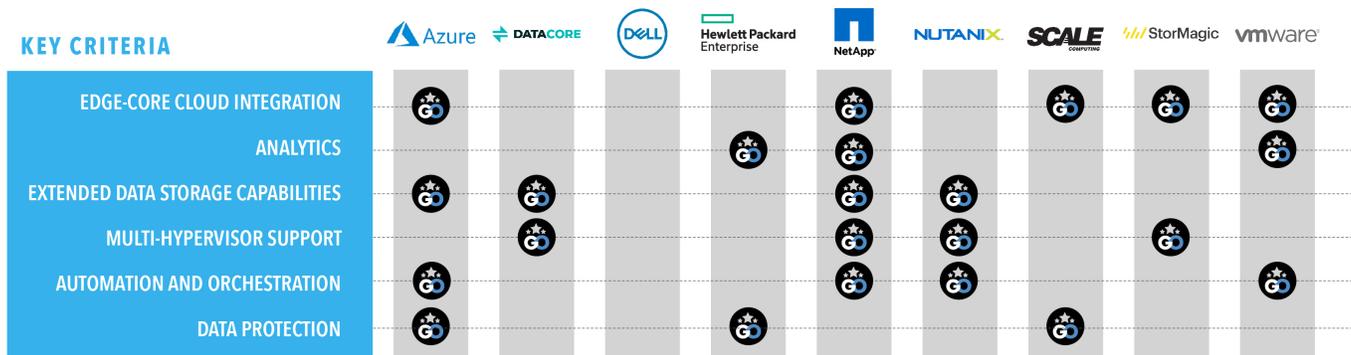
Organizations can select their preferred licensing model but more and more are switching to subscriptions to move CAPEX to OPEX and getting more cloud-like resource consumption experience based on pay-as-you-go mechanisms, while partially avoiding over-provisioning. In fact, some vendors started to take this model to the next level by providing the entire solution as-a-service including hardware, but still installed on the premises of the user.

5. Key Criteria

Not all of the organizations experience the same evolution regarding their IT infrastructure needs in terms of data growth, regulations, security, or changing business requirements, but on the other hand, it is important to choose infrastructures that are able to keep up with the adoption of new technologies and be prepared for potential requests coming from the business. Having an IT infrastructure that can stay efficient, agile, and aligned with the latest innovations brings several benefits to the entire organization, which can start new initiatives quickly and improve its competitiveness when needed, without a significant impact on the budget. HCI dramatically simplifies the IT infrastructure stack but this could create limitations in terms of flexibility at times. Choosing the right solution will limit this risk to areas that are less critical for the organization. The figure below outlines the noteworthy vendors for each key criteria.



GO Award Winners for HCI Capabilities, Innovation, and Key Metrics



Edge-Core-Cloud Integration

With more organizations considering hybrid cloud strategies, the integration between the hyper-converged infrastructure and the cloud is now critical. At the same time, edge computing use cases are becoming more common as well, forcing organizations of all sizes to look for specific solutions in this area that are simple to manage at scale and with a low TCO. In fact, one-size-fits-all HCI doesn't exist; therefore, it is very hard to find a solution that checks all the boxes when it is time to combine edge, core, and cloud functionalities.

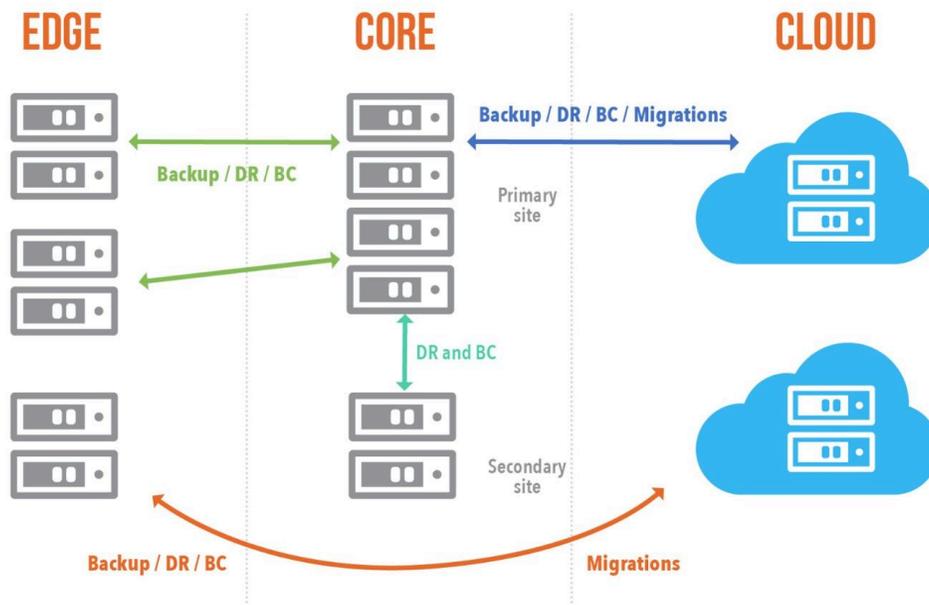
HCI solutions for edge computing have some peculiar characteristics that are quite difficult to find in other types of solutions:

- **Stronger security and data encryption.** Edge appliances are installed in remote locations with limited security; data stored in them should be protected against theft or any other risk associated with unsecured environments.

- **Efficient with limited resources.** Usually at the edge, there are few VMs to manage and they usually perform basic tasks. It is pretty common, in fact, to find very small installations starting with two nodes based on very limited hardware. This is a type of configuration that not all HCIs support efficiently.
- **Backup and remote replication.** One of the most important features for HCI at the edge is to have embedded backup and remote replication capabilities. In fact, in most of the cases, these installations are unmanned and data protection policies are defined on the central management system placed in the core (enterprise data center) or cloud.
- **Cost-effective.** Another characteristic of this type of system is the overall cost of the solution. This is why many edge-focused HCI systems give the option to use KVM as a hypervisor.

Core and cloud are natural extensions of HCI at the edge. They provide a secure backup repository and a remote replication target from which it is possible to quickly recover and run VMs in case of a failure in the edge site (See Figure 5). The same goes for solutions focused on SMB: the cloud is a perfect match and for similar reasons (backup and DR), mostly because the organization needs advanced data protection and wants to provide business continuity without the cost and complexity of managing multiple sites.

Figure 5. Examples of HCI edge-core-cloud integrations



Source: Gigaom (2019)

In large enterprises, the integration between on-premises hyper-converged infrastructures and the cloud goes beyond backup, disaster recovery, and business continuity thanks to additional services offered by the hypervisor or the HCI vendor aimed at improving VM and data mobility, synchronization,

and migration between platforms. Common cloud integrations offered by HCI vendors include:

- **Virtual appliance:** The HCI vendor offers a version of its platform that can run on the infrastructure of one or more cloud providers. This is a very simple and effective solution for all those enterprises that want to keep costs down while keeping data and applications synchronized to a remote location and ready to be started in case of a disaster. Usually available on the cloud provider's marketplace or as a software package to install in a VM instance of the cloud, it can be managed by the same console that oversees the rest of the cluster and the feature set is similar to what is available on the on-prem version of the product, further simplifying adoption and management.
- **Complete environment:** Some vendors are able to offer full-featured virtual managed infrastructures that run natively on cloud infrastructures (VMware on AWS or Microsoft Azure for example). This solution is more efficient than the previous one and aimed at operations that need to constantly run some of their workloads on the cloud. Integration between the on-premises installation and the cloud can happen in different ways: with tools offered by the hypervisor, by installing virtual appliances alongside the cloud infrastructure, or third-party tools specialized in data mobility. This type of solution allows support of large edge infrastructures without needing to make further investments in existing data centers, as a backend for IoT applications, or for all those use cases where there is the need for cloud-native services alongside legacy applications connected to remote environments.
- **Cloud services:** Some HCI vendors offer the ability to take advantage of services available natively on the cloud for operations like snapshot offload, data protection, backup, and so on. Also, in this case, a virtual appliance or a complete environment could be available to simplify re-hydration operations, migrations, or to just run VMs that can access data previously stored in the cloud. By taking advantage of S3 (object storage) repositories, for example, the end user can make copies of data in the cloud for data protection but also create copies of data that can be reused by other applications; for example, analytics or other use cases including development and test scenarios.

All the described solutions for integrating HCI with the cloud are targeted to different use cases and the overlap is limited. In fact, some HCI vendors offer all of them to give their customers more freedom of choice regarding the architecture design of their entire infrastructure. Alongside what is usually offered by the single HCI vendor, commercial hypervisors offer a broad ecosystem of third-party solutions that can support this type of migration or complement the existing ones.

Noteworthy Implementations

VMware HCI

VMware HCI stack is now available on Amazon AWS, Microsoft Azure, and IBM cloud, enabling organizations to have the same user experience they have for their on-premises HCI installations on the cloud. VMware Cloud on AWS is also the foundation of VMware Site Recovery Service, a cost-effective pay-as-you-go solution for disaster recovery of VMware environments. VMware ESXi customers can take advantage of virtual appliances provided by their HCI vendors and attach them to one of the cloud providers mentioned above for seamless data synchronization and migration to the

cloud within a familiar environment.

Scale Computing HC3

HC3 Cloud Unity has a simple and neat solution developed with Google Cloud Platform(GCP), which is a perfect match for its edge and enterprise platforms. Scale Computing customers can run the same identical environment they have on their premises on the cloud, using it for remote backup and disaster recovery, or for hybrid cloud use cases. The user can manage the entire environment from a single console, enabling resource movement from on-premises clusters to the cloud and back.

Microsoft Azure Stack HCI

Thanks to the recently introduced Azure Stack HCI, now Microsoft has a very comprehensive set of solutions that enable organizations to deploy and manage legacy VMs (Azure Stack HCI) and cloud-native applications (Azure Stack) on-premises seamlessly integrated with cloud services provided by Microsoft Azure. Azure Stack HCI can also be deployed in cost-effective 2-node configurations for edge use cases while maintaining the same feature set and capabilities of larger configurations.

StorMagic SvSAN

StorMagic SvSAN has proven over the years to be an excellent solution for edge use cases and now, thanks to direct support for KVM and additional encryption features, the company is able to provide a very cost-effective, end-to-end, and secure solution for large-scale edge deployments. The ability to run on very small 2-node configurations while granting good high availability associated with centralized management, allows these solutions to keep costs at bay while providing high SLA standards in remote locations.

NetApp HCI

Even though NetApp HCI doesn't primarily focus on edge use cases, this solution has begun to see some deployments at the edge. NetApp HCI takes advantage of all the integrations made possible by NetApp Data Fabric, a set of data services designed to provide consistent capabilities across public cloud and on-premises storage systems. NetApp HCI can easily move data to other on-premises NetApp storage systems as well as to public cloud providers through standard protocols like S3, or by replicating data through NetApp SnapMirror and other tools.

Analytics

Analytics is now a fundamental component of any modern IT infrastructure. With the growing infrastructure complexity and with applications spreading between on-prem and cloud, it is becoming more and more important to have tools aimed at giving insights quickly about what is really happening, enabling the administrator to remediate potential issues before they actually happen. Virtualized and hyper-converged infrastructures are no exception and fortunately, there are several hooks from which

analytics tools can gather information and create logs to provide a better picture of what is happening and how to proceed, in order to get the most out of the infrastructure.

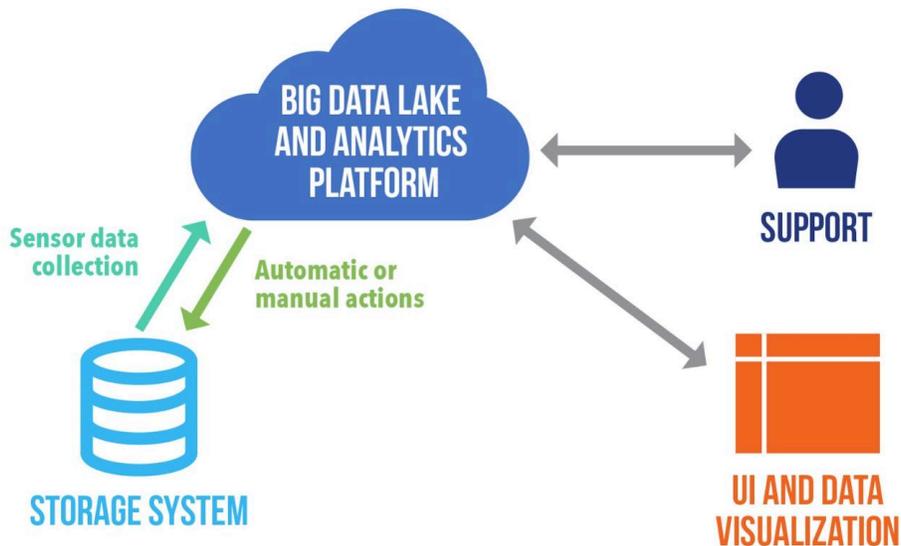
Depending on the hyper-converged platform, there are two different approaches to analytics for HCI:

1. **The HCI vendor provides the necessary tools** to collect, analyze, and visualize data on the entire stack. This is the best method to monitor the infrastructure and its behavior, no matter which hypervisor or VMs are running on top of it.
2. **The HCI platform can be integrated with external tools**, usually at the hypervisor level. This gives better visibility on applications and all upper layers including VMs but is less effective on hardware components of the infrastructure.

The best option would be a combination of the two that can monitor the entire stack but if not well implemented, this could easily result in additional costs and redundant information that can create confusion and delays in operations.

The main benefits of analytics for HCI are the same as what it is possible to get from similar solutions for other parts of the infrastructure, especially when the analytics platform collects data from all the installations worldwide and can make comparisons between a single system and ranges or average of single metrics (see Figure 6).

Figure 6. Logic Scheme of Modern Storage Analytics Platform



Source: Gigaom (2019)

These benefits include:

- **Faster troubleshooting.** Continuous data collection and visualization help operators have a complete view of the system and intercept issues before they actually become critical.
- **Better vendor support.** System data collected from every single system in the field is aggregated in a single big data lake and used to identify error patterns or system misconfigurations that will help support teams intervene quickly and preemptively to minimize service disruptions.
- **Improved capacity planning.** By comparing local with global trends about system resources allocations on similar systems and workloads, it is possible to make accurate forecasts about which and when new system resources will be needed to expand the system.
- **Continuously tuning and optimization.** Every single system can be compared to others or to best practice documents with the goal to find misconfigurations and suggest optimizations that can improve system health and performance.

Noteworthy Implementations

HPE InfoSight

HPE InfoSight, an industry-leading infrastructure analytics platform, is now available for HPE SimpliVity HCI and HPE Nimble dHCI systems, bringing the same functionalities available to HPE products to HCI as well. The solution provides the same level of detail available for Nimble Storage systems, including full-stack performance analysis up to the VM layer.

NetApp Active IQ

Active IQ is a complete and sophisticated analytics tool available across many NetApp products, which takes advantage of AI techniques to constantly improve pattern recognition and advise the end user about their HCI and storage infrastructure including troubleshooting and remediation, capacity planning, automated support, and so on.

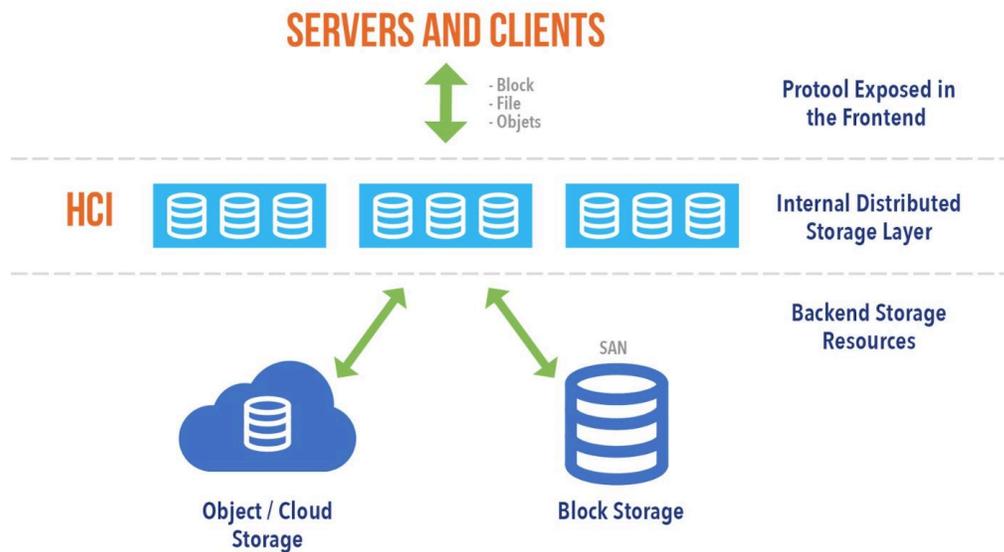
VMware vRealize Operations

VMware, thanks to the introduction of vRealize Operations 7.5, now has a very powerful tool to help sysadmins discover anomalies, work on remediations, and optimize configuration and performance of the entire stack. The product is available to all VMware customers, but now it has specific functionalities aimed at improving the performance of vSAN-based HCI environments.

Extended Data Storage Capabilities

With hyper-converged infrastructures serving more and more workloads, data storage for VMs is not enough to provide the storage resources needed to efficiently cover all use cases. Furthermore, to get the best flexibility and TCO, the HCI platform should be able to access external storage resources, locally and in the cloud, to offload cold data for backup and DR, and to expand the internal storage with existing resources for added flexibility (see Figure 7).

Figure 7. Logic Scheme of extended storage capabilities for HCI



Source: Gigaom (2019)

The main benefits of exposing file, block, and object storage services to external clients directly from the HCI platform is to keep the infrastructure simple and easy to manage while providing all the necessary access protocols to serve legacy applications that can't be virtualized or that need file and object interface. At the same time, having the ability to take advantage of existing storage resources and see them as part of the storage pool available to VMs and other services is of great help to smooth the migration path from legacy to hyper-converged infrastructure while minimizing the impact on the budget.

Storage exposed in the front-end by the HCI platform can include:

- **Block devices.** Data volumes usually presented to external servers via iSCSI. These volumes keep all the characteristics that are available internally to VMs and allow external servers to have access to shared storage resources without the need for further investments on a legacy SAN (storage area networking).

- **File Services.** Shared file systems exposed via NFS and/or SMB protocols. This type of service allows consolidation of file servers and network attached storage (NAS) systems on the HCI, removing external appliances and simplifying the infrastructure.
- **Object Storage.** \$/GB for object stores is very low and they are very scalable too. Delivering this service starting from a storage layer designed for primary workloads, such as the case of the distributed storage layer of an HCI, is really expensive and counterintuitive. But at the same time, the number of workloads that take advantage of object stores is growing daily. In fact, an increasing number of enterprise applications use it as a target to store their data and in most of the cases, it is relatively small capacities and active data sets that also require performance.

In all the cases described above, HCI can help with the consolidation and simplification of the IT infrastructure with important benefits both in terms of ROI for the HCI and an improved overall TCO. To take full advantage of these functionalities, the HCI platform should support nodes of different capacities and performance in the same cluster with the possibility of associating the right kind of storage media to every single service, helping to find the right \$/GB and performance figures for all types of workloads and data types. Strong multi-tenancy features, like QoS, should also be in place to prevent conflicts between different workloads and applications.

Some HCI solutions can also benefit from backend storage resources:

- **Block devices.** Taking advantage of an existing SAN to expand HCI storage allows the organization to reuse existing infrastructure resources, helping to limit the initial investment in the HCI platform while smoothing the migration process. This is particularly important when the depreciation cycle of an existing SAN is not concluded, allowing a smoother transition to HCI from a financial perspective.
- **Object storage.** Offloading data to external object stores, no matter whether they are in the cloud or on-premises, can save capacity on the primary storage tier and make more resources available to VMs in production. This helps the enterprises implement a two-tier strategy and improve the \$/GB of the infrastructure while delaying investments in primary storage.

Again, the benefits brought by extended storage functionalities include better flexibility of the platform, which leads to a smoother migration path and better overall infrastructure TCO, thanks to the ability to reuse existing infrastructure components. In the long term, this will allow more consolidation of workloads and data on a single HCI platform, which is easier to manage and expand over time.

Noteworthy Implementations

Datacore Hyper-converged Virtual SAN

Hyper-converged Virtual SAN is based on the same technology of SANSymphony and inherits all its characteristics, including the ability to virtualize existing storage resources and present them to VMs or other systems in the storage network. This allows end users to reuse existing storage systems and

smooth the migration path to HCI while optimizing the investment.

Nutanix Storage Services

Nutanix offers a complete set of storage services, including scalable file SMB/NFS volumes, S3-compatible object stores, and iSCSI targets accessible to external servers as well. All services are provided within the standard UI and management tools. This, associated with the ability of the Nutanix platform to accommodate nodes of different capacity and performance in the same cluster, allowing the consolidation of several workloads on the same cluster at reasonable costs.

NetApp HCI

HCI and SolidFire all flash arrays share a common set of control plane services with their Element OS, including the ability to dynamically provision resources and services and creating shared block volumes with external systems. The HCI system can replicate data to SolidFire and other ONTAP-based NetApp systems, as well as offload snapshots to S3 repositories. In addition, file and object storage services are available on the platform by installing NetApp Cloud Volumes ONTAP and StorageGRID on the HCI system, and they can take advantage of QoS functionalities of the system to separate workloads at best.

Microsoft Azure Stack HCI

Azure Stack HCI is based on Windows Server 2019 and Storage Spaces Direct, a powerful storage management layer that allows the creation of data volumes spanning across several server nodes with consistent data protection, which can be used for internal VMs as well as shared via standard block (iSCSI) and file protocols (SMB and NFS).

Multi-hypervisor Support

Even though many organizations, especially in the small/medium enterprise (SME) segment of the market, prefer to work with a single hypervisor to simplify management operations. At the same time, supporting multiple hypervisors on the same HCI platform can have dramatic benefits in terms of cost savings and efficiency in specific use cases or large environments.

VMware ESXi has the lion share in the hypervisor market but there are reasons why customers prefer other platforms or keep more than one in production. For example, as previously mentioned, KVM has the right characteristics for edge use cases, but it is not as good for traditional enterprise use cases. If the HCI platform supports both of them, the enterprise can keep the same HCI solution while switching the hypervisor depending on the needs of the single project with a minimum impact on operations. Another example can come from Hyper-V, which has been seeing a certain growth lately, mostly because of the success of the Microsoft Azure cloud and the integration possibilities between the two.

In many cases, multi-hypervisor support is all about freedom of choice, especially considering the

possibility that the IT infrastructure will expand to support new use cases in the future. On the other hand, supporting more hypervisors also has other benefits:

- KVM support usually comes as part of an end-to-end solution from the HCI vendor. Instead of supporting several KVM flavors, the HCI vendor concentrates its efforts on a single Linux distribution and packages it with its own software. The customer gets better prices and support with additional features designed for specific use cases that are not always available on commercial hypervisors.
- Having multiple choices at the hypervisor level helps to avoid lock-ins. Even though VMWare ESXi can't be considered a lock-in per se, company strategy could be influenced by several factors that will reflect on the IT infrastructure, including cloud, IoT, and edge computing projects. Having a second, cost-effective, solution based on the same HCI platform but with a different hypervisor can be of great help to keep costs down without disrupting processes in place.

In order to get multi-hypervisor support properly, the HCI vendor should provide all the migration tools between the two virtualization systems as well as an integrated console to hide the complexity of managing two platforms concurrently.

Noteworthy implementations

Nutanix Acropolis

Nutanix supports commercial hypervisors (VMware ESXi and Microsoft Hyper-V) alongside its KVM-based AHV hypervisor and tools to migrate VMs across different environments. This gives enterprises plenty of choices when it is time to select the right compromise between features and cost for each single VM.

NetApp HCI

NetApp offers multi-hypervisor compatibility, including VMware ESXi and RedHat Virtualization Platforms alongside OpenStack, NetApp Kubernetes services, and RedHat OpenShift, thereby, offering large enterprises and ISPs several options for deployments in complex environments that are beginning to place virtualization and container orchestration side by side in the same infrastructure.

DataCore Hyper-converged Virtual SAN

DataCore offers compatibility to ESXi, Hyper-V, Linux KVM, and Citrix Xen Server, giving enterprises one of the most comprehensive lists of hypervisors from which they can freely choose what is best for their needs.

Stormagic SvSAN

SvSAN is a cost-effective solution that finds most of its success in edge infrastructures. It is compatible with VMware ESXi, Microsoft Hyper-V and open-source KVM hypervisors. SvSAN offers an end-to-end solution that includes advanced encryption functionalities to improve remote site security and better control over the infrastructure.

Automation and Orchestration

Even though initial system deployment is automated and simplified for all HCI systems, API and automation tools are now fundamental components of any large IT infrastructure in day-to-day operations. With applications now divided into several components that can scale differently depending on momentary needs, and business that requires faster and faster resource provisioning, automation and orchestration become essential tools for administration and operational efficiency.

The two major use cases for which automation and orchestration are of great help are:

- **Fast provisioning.** Virtualized environments have several advantages in terms of resource provisioning when compared with physical systems, and with hyper-convergence, this process is even improved. In fact, the tightly integrated components of an HCI solution allows for greater visibility on the entire stack, making the whole process easier and controllable.
- **Self-service provisioning.** Many organizations want to build a cloud-like experience for development teams and business units with the goal to simplify and speed up provisioning processes. The infrastructure should operate quickly upon requests usually made through self-service portals available to end users.

Organizations should always be aware of the possibilities offered by its HCI regarding API compatibility and command line interface (CLI) features, and should always ask for integrations with leading orchestration and configuration tools available in the market, or see if the HCI vendor provides them. Another important aspect to take into account is that many infrastructures are usually configured and maintained with tools like Ansible, Puppet, Chef, or Terraform, therefore the HCI vendor must be able to provide and manage resources and services through these tools as well.

Noteworthy implementations

VMware vRealize Suite

VMware vRealize Suite is available to all VMware customers, including vSAN as well as any other HCI organization. It provides a broad spectrum of products and functionalities to simplify monitoring, automation, and orchestration of operations in VMware and hybrid-cloud environments. As it happens for many VMware-based products, vRealize can count on a large active user community and support from a long list of third-party vendors.

Nutanix Calm

Nutanix Calm is an orchestration platform that allows Nutanix customers to perform repetitive and complex operations easily in complex and hybrid-cloud environments. It also includes the possibility of publishing self-service portals that internal teams can use to deploy complete environments in minutes while keeping control over security and processes.

NetApp HCI

SolidFire, and now its HCI platform, are well known in the industry for their API and the work done with the open-source community to support tools and plug-ins for most automation and configuration systems available. At the same time, the SolidFire team also contributed to the creation of several interfaces, such as OpenStack Cinder, Kubernetes Container Storage Interface (CSI), and recently Trident, for persistent storage container orchestration.

Microsoft System Center

System Center is a powerful suite of system management products that include Operations Manager and Orchestrator. For large organizations adopting Azure Stack HCI, System Center is the right tool to simplify monitoring and automation of the entire stack across on-premises infrastructure and Azure cloud.

Data Protection Services

The primary goal of hyper-convergence is infrastructure simplification. In this context, integrating data protection services in the HCI solution allows it to simplify processes and operations, contributing to improving the overall TCO of the infrastructure. At the same time, if the HCI solution has the right tools in place, DR orchestration can also be extremely simplified.

Most solutions in the market offer integrations with third-party data protection software at the hypervisor level but some integrate backup and DR directly onto their platform. Both approaches have advantages and disadvantages and depend on user needs.

- **Hypervisor-level data protection** allows common, external, data protection backup platforms to serve heterogeneous environments and, potentially, move data across them when necessary. This is the most common solution among enterprises with a variety of systems to support.
- **Native HCI data protection** is optimized for the specific HCI solution and usually more efficient than hypervisor-level data protection, allowing fast data retrieval, providing more insights on protected data, and enabling easier DR orchestrations. The drawback is that it is limited in supporting complex and legacy environments; most of its installations are in small and mid-sized organizations.

HCI vendors usually leverage snapshots for data protection and remote replication. Advanced schedulers, with a simple-to-use UI, offer control over backup, retention policies, and data placement.

Snapshots are space-efficient, so they can be scheduled frequently to improve the recovery point objective (RPO) and recovery time objective (RTO), which can be a few minutes. Usually, these tools are not as sophisticated as external data protection solutions and some lack the level of granularity of file-level restore inside the single VM or the agents to provide backup at the application level (as in the case of databases, for example). Integrating a third-party data protection solution with the tools provided by the HCI platform usually provides the most effective combination of functionality for large scale environments while keeping the system easy to use and manage.

Cloud object stores or secondary HCI systems are common targets for backups and they can also be leveraged to implement DR strategies with the DR orchestrator tools provided by the HCI vendor.

Noteworthy Implementations

HPE SimpliVity

HPE SimpliVity has always prioritized data protection tools for backup and DR orchestration as the core of its solution. With strong deduplication, SimpliVity is able to make local and remote copies of VMs alongside incremental snapshots that take up a small footprint. HPE SimpliVity provides complete quick restore function, multi-site support, advanced scheduling, and tools to manage failover and failback in case of disasters.

Scale Computing HC3

Scale Computing has an easy-to-use solution for backup and DR of single and distributed edge-cloud infrastructures, allowing administrators to take snapshots of VMs in remote locations and centralize them on a central repository, in the cloud, or in a core datacenter cluster, that can also be used for DR purposes.

Microsoft Azure Stack HCI

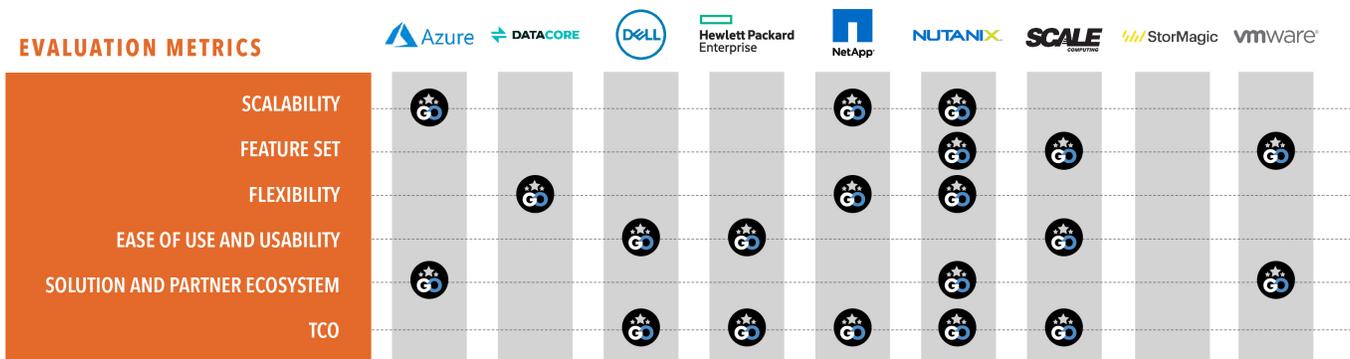
Azure Stack HCI includes data protection and DR tools seamlessly integrated with Azure Cloud. The end user can access these services directly from the management console and enable data protection with a few clicks. This functionality is also particularly interesting for enterprises with remote sites that want to simplify backup infrastructure and operations while taking advantage of cloud services.

6. Critical Features: Impact Analysis

As described earlier in the introduction, this report analyzes the impact of critical features of HCI solutions available in the market and puts them in context with the evaluation metrics that are usually at the core of strategic decisions. These strategic decisions can then be based on the impact each evaluation metric can have on the infrastructure, system management, and IT processes already in place with particular emphasis on ROI and TCO. The figure below outlines the noteworthy vendors for each evaluation metric.



GO Award Winners for HCI Capabilities, Innovation, and Key Metrics



The table below (see Table 1) reports a weight on a scale from 1 to 5 for each feature relative to the metrics analyzed, but every enterprise should check these values against the needs of his or her organization. For example, if flexibility is not crucial for an organization, multi-hypervisor support will have a smaller impact not only on flexibility but also on the rest of the metrics, including TCO or feature set.

The following table provides criteria to help organizations make informed decisions when evaluating a new hyper-converged infrastructure system. It is also important to note that the information reported here is relative to common enterprise scenarios and should always be mediated with the needs of the specific enterprise organization, local regulations, and other environmental constraints. In fact, the impacts of these features are not the same for everyone.

Table 1. Critical Features Impact on Metrics

	SCALABILITY	FEATURES SET	EASE OF USE/USABILITY	FLEXIBILITY	SOLUTION / PARTNER ECOSYSTEM	TCO
EDGE-CORE - CLOUD	2	3	3	5	1	4
ANALYTICS	2	2	5	1	1	5
EXTENDED STORAGE	5	4	3	4	1	3
MULTI HV	3	5	1	4	4	3
AUTOMATIC/ ORCHESTRATION	4	4	4	3	3	3
DATA PROTECTION	1	3	4	3	5	3

Scale: 1-lowest | 5-highest

Source: Gigaom (2019)

Best Impact on Scalability

Absolute scalability is not merely the number of nodes that an HCI solution can support in a single cluster. This alone does not say enough about a solution that aims at replacing most of the data center infrastructure. When the goal is to build a modern infrastructure that can consolidate several types of workloads, applications, and access methods on the same system, true scalability also means the ability to mix nodes of any type and size, as well as support multiple protocols concurrently.

Nutanix Acropolis

Nutanix has a very flexible architecture design that allows the consolidation of a large number of primary and secondary workloads, applications, and data in a single cluster while harmonizing overall system management, providing extended analytics and tools to simplify operations and many IT processes.

NetApp HCI

NetApp HCI, based on Solidfire technology, has already proven its scalability with important deployments in large enterprise organizations and ISPs. With the ability to leverage other software-defined storage (SDS) products available in the NetApp product suite, NetApp HCI can consolidate several types of data in a single system while granting the necessary multi-tenancy features needed to provide consistent and predictable performance.

Microsoft Azure Stack HCI

Azure Stack HCI is a very flexible solution, starting with a minimum of two nodes and capable of growing to large multi-PB clusters. Even more so, thanks to Storage Spaces Direct, Azure Stack HCI is able to consolidate several workloads and data on the same system, taking advantage of several types of devices including hard disks, NVMe-based flash memory, and Intel Optane.

Best impact on Feature Set

Comparing the HCI solutions, the most interesting ones for enterprises are those that provide the largest feature set, which translates into more use cases and options for infrastructure consolidation. This is also true for solutions designed for small and medium businesses and edge computing but due to the size of infrastructures and IT teams, in these cases it is more about features aimed at simplifying operations and reducing both total cost of acquisition (TCA) and ownership (TCO).

Nutanix Acropolis

Nutanix has one of the broadest feature sets in the market, enabling its customers to consolidate practically every kind of workload on its platform. Now, thanks to the introduction of Mine for

secondary data management, and the growing number of tools for cloud integration, Nutanix has the ability to offer one of the most comprehensive feature sets in the market.

VMware

VMware, thanks to its leadership position in infrastructure virtualization, has developed the most comprehensive end-to-end stack, offering the best hypervisor and the tools to manage it at any scale and across different environments, including public cloud. In addition, its HCI solution is tightly integrated with the hypervisor for the best performance and seamless migration experience from traditional environments for sysadmins.

Scale Computing HC3

Scale Computing has built one of the most complete end-to-end solutions for SMB, distributed enterprise, and edge computing use cases. Easy to use, cost-effective, and with an all-included licensing model that includes multi-site management, backup, DR and migration tools. At the same time, Scale Computing is also available on Google cloud, giving its end users the possibility to expand the infrastructure on the cloud.

Best Impact on Flexibility

Flexibility is one of the key characteristics of HCI solutions, enabling organizations to respond quickly to ever-changing business requirements. Factors that make this possible are: ease of cluster expansion; multi-hypervisor support for the best control over cost and functionalities; additional functionalities and tools that allow the consolidation of more workloads on the same infrastructure; and tools to expand the infrastructure to the cloud.

NetApp HCI

NetApp HCI is a very flexible enterprise solution thanks to the ability to expand compute and storage independently. At the same time, it is based on the same software that runs on Solidfire arrays, making it possible to build hybrid infrastructures for complex environments. In addition, NetApp HCI can take full advantage of the NetApp Data Fabric solution for data mobility and cloud management.

Nutanix Acropolis

Nutanix offers a well-designed and highly-configurable architecture integrated with a set of additional tools that enables administrators to deploy and expand their infrastructure quickly according to business needs. Integration with the cloud and advanced data services add further flexibility to the product, including disaster recovery orchestration and optimizations around data placement, particularly useful for the consolidation of primary as well as secondary workloads in the same system.

Datacore Hyper-converged Virtual SAN

Datacore Hyper-converged Virtual SAN brings the flexibility of storage virtualization to hyper-convergence. With a very user-friendly and capacity-based licensing model, organizations can take advantage of existing storage resources and make them available to the hyper-converged infrastructure, but the systems also work in the opposite way, giving organizations plenty of opportunities when it is time to expand the cluster or decommission old storage systems. At the same time, the storage layer can be expanded independently from the compute nodes, improving granularity and efficiency.

Best Impact on Ease of Use and Usability

Ease of use and in general, the usability of features offered by the hyper-converged infrastructure, is of great importance for two reasons. On one hand, an experienced sysadmin can manage a larger infrastructure while on the other hand, the system can be managed by a less skilled operator. This is of help for those organizations trying to simplify their IT operations and standardize processes and roles as much as they can. Additionally, many small and medium organizations usually have limited resources in their IT departments, and people have to cover different roles. Easy-to-use, analytics-backed HCI systems, well-integrated with the upper layers, can play an important role in the simplification of the infrastructure and its usability.

HPE SimpliVity and dHCI based on Nimble Storage

HPE platforms can take advantage of the award-winning Infosight analytics to simplify most system administration, troubleshooting, support, optimization, and capacity planning operations. Infosight leverages big data and machine learning algorithms to compare every system in the field with all the other installations to discover anomalies, analyze trends, and give the administrator valuable advice about potential issues and remediations before they reach a critical state.

Dell EMC VxRail

Dell offers all VMware HCI functionalities and a very simple, familiar feeling to every sysadmin that has already experienced virtualization, without the hassle of building and integrating the single components. In fact, it is completely integrated and offers a broad set of solutions to help enterprises with deployment automation, system management, and upgrade processes.

Scale Computing HC3

HC3 is easy to use and manage, with one of the best support services experienced by enterprises. Designed to be run by small teams with limited time resources, the management console is very intuitive and integrated within the cluster, helping inexperienced operators to keep the system always efficient.

Best Impact on Solution and Partner Ecosystem

Solution and partner ecosystem is crucial for HCI enterprise adoption and sustainability over time. Even though most of the integrations are made at the hypervisor level, direct integration and support from third parties are of great help in building integrated infrastructures, accessing best practices and simplifying system management activities, contributing to a better infrastructure TCO and user experience.

VMware HCI

VMware has always offered the possibility to integrate its products with third-party solutions through APIs and plug-ins, with a positive attitude towards technology partnerships. This, alongside its market leadership positions, contributed to building an impressive ecosystem around its entire stack, including its HCI solutions, both from software and hardware perspectives. In fact, even though VMware is a software-only solution based on standard HW that enterprises can select from a HW compatibility list, there are other options including a pre-certified system, or solutions like Dell EMC VxRail which offer end-to-end stacks including hardware and software with all the necessary tools and automation to further simplify initial deployments and upgrades of the systems.

Nutanix Acropolis

Nutanix has developed a solid ecosystem around its product portfolio and now it has direct support from many primary hardware and software vendors. This strategy enables its customers to pick up the best solution for their needs and build efficient infrastructure stacks that can take full advantage of Nutanix's characteristics and capabilities.

Microsoft Azure Stack HCI

Microsoft, thanks to the introduction of Azure Stack HCI and its tight integration with Azure Cloud, can provide a solid ecosystem of solutions that spans between on-premises, edge, and cloud. Even though the number of third parties working with Hyper-V is limited compared to the market leader, the ecosystem is growing pretty quickly, due to a brilliant partnership strategy and the success of its cloud solutions.

Best impact on TCO

TCO is always a difficult metric to evaluate. It includes TCA and differs from organization to organization due to a lot of variables. To find the solutions that have the best impact on TCO, we take into consideration all the metrics assessed in the previous chapters, giving an overall evaluation for the product in general terms.

Nutanix Acropolis

Nutanix has been able to build one of the most compelling solutions in the market. Even when it doesn't have the best feature in a category, its products always score very high and are solid. Customers appreciate the ease of use, support, and the enviable partner ecosystem.

Dell EMC VxRail

VxRail adds additional benefits on top of VMware HCI stack, further simplifying procurement, first deployment, and system management of a VMware-based virtualized infrastructure. This out-of-the-box solution is very popular among enterprise end-users for its solidity and simplicity. Dell is expanding its offering with VMware Cloud on Dell EMC so that large organizations can take advantage of these products on-premises yet sold through a subscription model and managed by Dell.

HPE SimpliVity

HPE provides a complete out-of-the-box solution, a capillary partner network, and a very easy-to-use system that perfectly matches the needs of many mid-size organizations and enterprises. Support for multi-HV also gives freedom of choice to its customers, while InfoSight still remains of great help to simplify system management.

NetApp HCI

NetApp HCI is a very compelling solution for large enterprises and ISPs, especially for those that are already NetApp customers. In general, the solution is well integrated with NetApp Data Fabric products and services, making data mobility across hybrid cloud environments easier and migration from other storage systems seamless. In addition, the extended hypervisor and container orchestration support make this solution really interesting for those organizations that are looking for workload mobility as well.

Scale Computing HC3

HC3 focuses on system simplicity with an iPhone-like approach regarding the user experience. In fact, by providing an end-to-end stack that includes the hypervisor, this solution is well received by organizations for both edge and SME use cases; and the proven quality of support attracts others. Even though the ecosystem is not as expansive as other competitors, they have been working on the ecosystem and made some progress. In addition, Scale Computing provides many features that are well suited for small, mid-sized, and distributed enterprises.

7. Near-Term Game-Changing Technology

Replacing traditional infrastructure, or a part of it with HCI, brings many benefits. Every organization should always be aware of what will come next so as to not be caught off guard and to be able to plan in advance regarding an infrastructure strategy. To do that, it is necessary to take a look at what is going to happen in the market and the potential evolution of technology that can improve efficiency, reduce friction, or simplify the infrastructure stack even further.



GO Award Winners for for Near-Term Game Changing Technology



Container Support

Kubernetes is quickly becoming the de-facto standard for container orchestration. More and more enterprises are moving their project from development and test phases to production, and all HCI vendors are working to support it in the best possible way. At the moment there are three different common approaches and practically all of the vendors are working on at least two of them in parallel:

- **Support for an existing Kubernetes distribution.** Many HCI vendors are working with market leaders, such as RedHat (OpenShift) or Google (Anthos), to provide certification and support for these platforms.
- **Provide a Kubernetes distribution.** Some HCI vendors have their own Kubernetes distribution optimized for their stack and with tools that are aimed at simplifying deployment and day-to-day administration tasks.
- **Generic support for Container Storage Interface (CSI) and Trident.** This allows organizations to be free to choose what best fits their needs without committing on distribution yet. This is the best option if they are still in an evaluation/test phase or if they chose an alternative distribution.

In any case, container support is becoming an important feature for many enterprises and this will become even more critical in the next year.

The most notable vendor in this space is **NetApp**, which has directly contributed to CSI development, Trident (persistent storage orchestration for containers), and offers the broadest solution, including support for its own Kubernetes distribution among others.

IT Composability (and NVMeOF)

All vendors have already started to integrate NVMe devices in their solution and this will be followed by the introduction of NVMeOF. Composable infrastructures allow users to organize hardware resources in pools and reconfigure the cluster as needed in a matter of minutes. NVMeOF, by giving the possibility to access remote storage devices and use them as if they were local to the server, allows composability with standard components.

Even though hyper-convergence is not the primary use case for composability, many vendors are working on new architectures that provide fast resource provisioning while decoupling storage from computing for better scalability and granularity.

The leader in this category at the moment is **HPE** with its Synergy product line, while many startups are developing interesting solutions around similar concepts.

Serverless Frameworks

With the success of serverless computing – where relatively small pieces of code can be triggered by events – enterprises are looking for technologies that do the same in their infrastructure, thus providing a public cloud experience in a private cloud. Some vendors are already working toward this direction, providing elements usually found in the public cloud on top of their HCI stack, and a serverless framework is a natural consequence of that. If well integrated with the rest of the infrastructure, a serverless framework could easily provide services both at the infrastructure level – to improve automation, for example – as well as to developers for their applications.

In this context, **Microsoft** Azure Stack is very well positioned to extend public cloud services to on-premises infrastructure for large organizations, while Azure Stack HCI provides services for VM-based workloads.

Advanced Security Features

With HCI infrastructures becoming more and more relevant in large organizations, security is quickly rising in the list of requested features by large organizations. Large hyper-converged infrastructures allow users to consolidate several workloads, concentrating not only data and VMs but also more sysadmin accesses while exposing services to a growing number of external applications.

in addition, with the expansion of the infrastructure to the cloud and in large distributed environments, having a tighter control on the entire infrastructure stack is fundamental in order to comply with demanding regulations and organizational policies. As a result, some vendors are working in some critical areas such as:

- Role-based access control, monitoring, and auditing;
- Data encryption at rest and key management;
- VM micro-segmentation and other techniques to improve security at the application or single workload level.

In this regard, **Nutanix** has already demonstrated a series of tools aimed at helping system and security administrators to simplify and improve many security procedures all across the infrastructure stack.

The following table (see Table 2) reports a weight on a scale from 1 to 5 for each feature relative to the metrics analyzed. The table summarizes the impact of the game-changing features on the metrics an enterprise needs to evaluate, with one having a lower impact and five having the most impact. Every enterprise should check these values against the needs of their organization. For example, if scalability is not crucial for an organization, NVMeOF and composability will have a smaller impact not only on scalability but also on the rest of the metrics, including flexibility.

Table 2. Game-Changing Impact Features on Metrics

	SCALABILITY	FEATURES SET	EASE OF USE / USABILITY	FLEXIBILITY	SOLUTION / PARTNER ECOSYSTEM	TCO
CONTAINER SUPPORT	3	4	1	5	4	2
COMPOSABILITY / NVMEOF	5	4	1	4	1	3
SERVERLESS	4	4	3	4	1	2
SECURITY	1	4	3	3	4	5

Scale: 1-lowest ▼ | ▲ 5-highest

Source: Gigaom (2019)

8. Conclusion

HCI simplifies infrastructure with a scale-out approach that partially eliminates the need for hardware sizing while adding better granularity and linearity in capacity and performance upgrades. However, given the range of features, the emphasis shifts to understanding the differences between implementations in order to choose the best solution for the organization. While HCI with standard, table-stakes features are mature, the market is very crowded with multiple HCI solutions offering key criteria and innovative features that can have a significant positive impact on organizations. Now more than ever, given the growing complexity of expanding infrastructure to the public cloud and the edge, it is crucial to select a solution that can last for a long time and respond quickly to ever-changing business requirements.

The HCI market share leaders are VMware and Dell EMC VxRail on one side and Nutanix on the other. VMware is the market leader in virtualized infrastructure and is expanding its reach to the hybrid cloud, thanks to a series of new products and services built in partnership with major service providers. Nutanix has been able to build a compelling alternative end-to-end solution that has a technical lead and is expanding to become more hybrid as well. Even when VMware (or VMware/Dell) and Nutanix do not come first in a single category, their stacks are always well positioned.

The runner-ups are HPE, Microsoft, and NetApp with very interesting solutions, each one of them with clear differentiators and market focus.

9. About Enrico Signoretti



Enrico has 25+ years of industry experience in technical product strategy and management roles. He has advised mid-market and large enterprises across numerous industries and software companies ranging from small ISVs to large providers.

Enrico is an internationally renowned visionary author, blogger, and speaker on the topic of data storage. He has tracked the changes in the storage industry as a Gigaom Research Analyst, Independent Analyst and contributor to the Register.

10. About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

11. Copyright

© [Knowingly, Inc.](#) 2019. "Key Criteria for Evaluating Enterprise HCI" is a trademark of [Knowingly, Inc.](#). For permission to reproduce this report, please contact sales@gigaom.com.