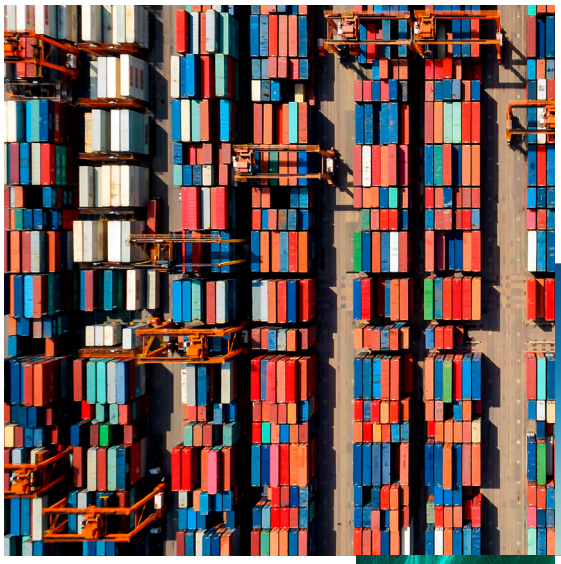

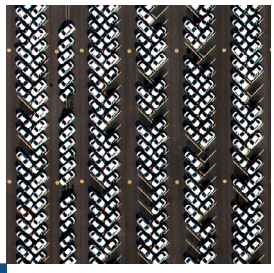

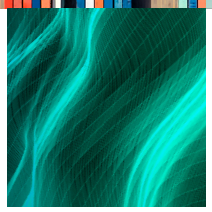

Not All Hyperconverged Infrastructure Solutions Are Created Equal

A Guide to Understanding Common HCI Architectures

NetApp

Table of Contents

Diagrams and tables in this document provided by IT Brand Pulse, 2020.

The Ripple Effect of the Cloud: One Change Has an Enormous Impact

The success of the cloud is driving dozens of industry initiatives, including hyperconverged infrastructure (HCI). HCI plays a key role in addressing these four ripple effects of the cloud.

Effect #1: Experience and expertise are needed for multiple types of on-premises and cloud data domains.

IT organizations must find a way to efficiently support hybrid clouds that consist of data in their on-premises private cloud and data in the public cloud.

Hybrid Cloud

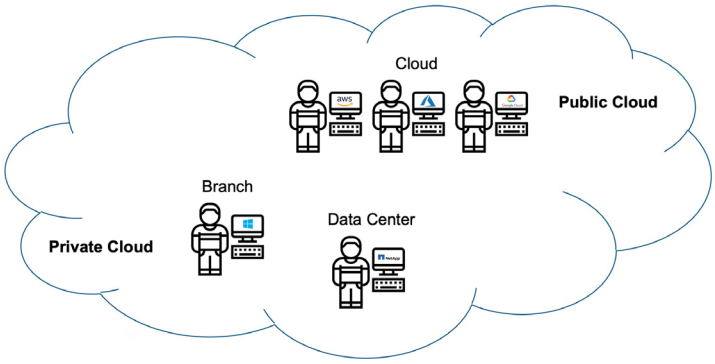

Effect #2: Disaggregation is needed to deliver performance efficiently and cost effectively to different workload-specific data domains.

IT organizations are challenged to meet performance SLAs for databases, virtualization, containerization, analytics, and artificial intelligence (AI) and machine learning (ML) workloads—all with significant differences in the compute and storage resources needed.

Hybrid Cloud

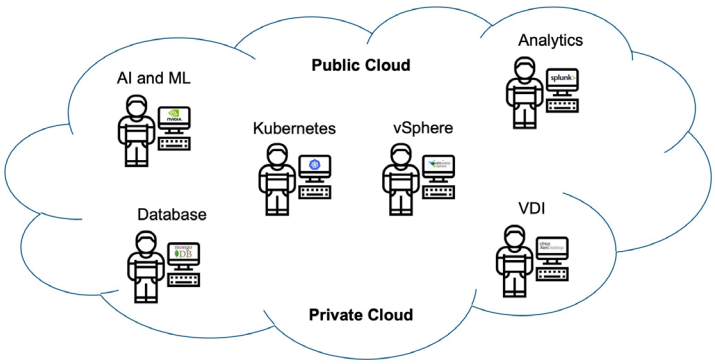

Effect #3: A Cloud Experience Is Preferred

In 2020, 14 years after the launch of cloud computing, the majority of IT professionals prefer a "cloud experience" that includes fast time to deployment, the ability to provision only what's needed, and granular scalability. A growing number of organizations also prefer a subscription for an on-premises infrastructure, where the service provider owns and operates the equipment behind the data fabric.

Effect #4: A Scalability Gap Emerges

The ripple effect of the cloud means that most organizations are in some phase of implementing initiatives for more efficient hybrid clouds, digital transformation, data analytics, and AI and ML. The need for resources to support these strategic initiatives is growing faster than IT budgets, leaving organizations looking for ways to make up the difference. HCI helps bridge that gap.

The Cloud Experience

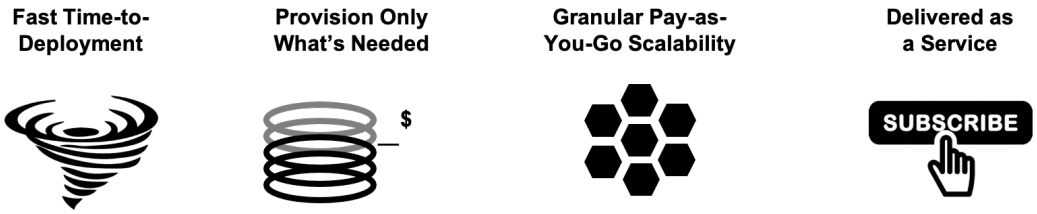

A Scalability Gap

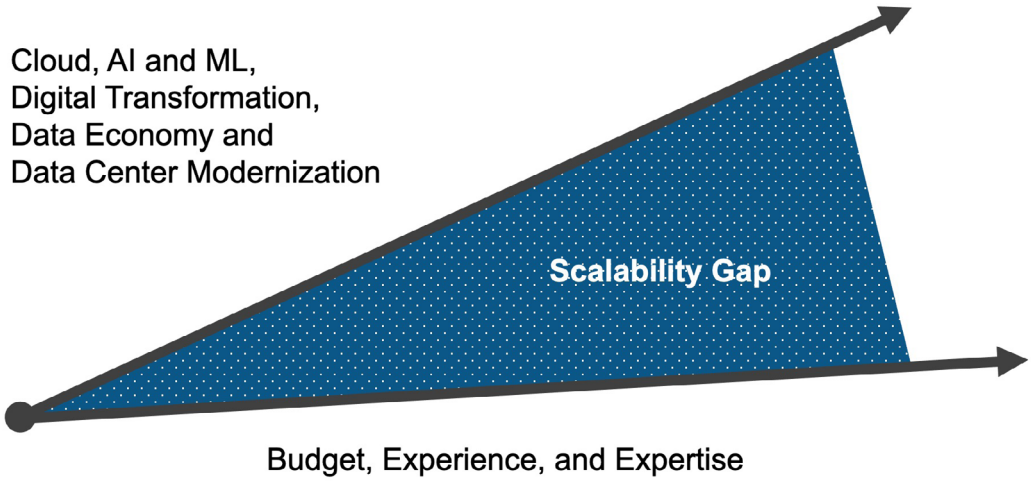

HCI Bridges the Scalability Gap

For decades, IT organizations have been researching, evaluating, purchasing, and integrating their data center infrastructure. The server and storage infrastructure of a typical data center is refreshed every 3 to 5 years. Up to a year of that time is invested in planning and deploying the new environment.

Simply put, HCI is data center infrastructure, preintegrated by vendors after careful design and testing. HCI bridges the scalability gap with proven templates that increase availability while simultaneously lowering time to deployment and cost.

Scaling with HCI

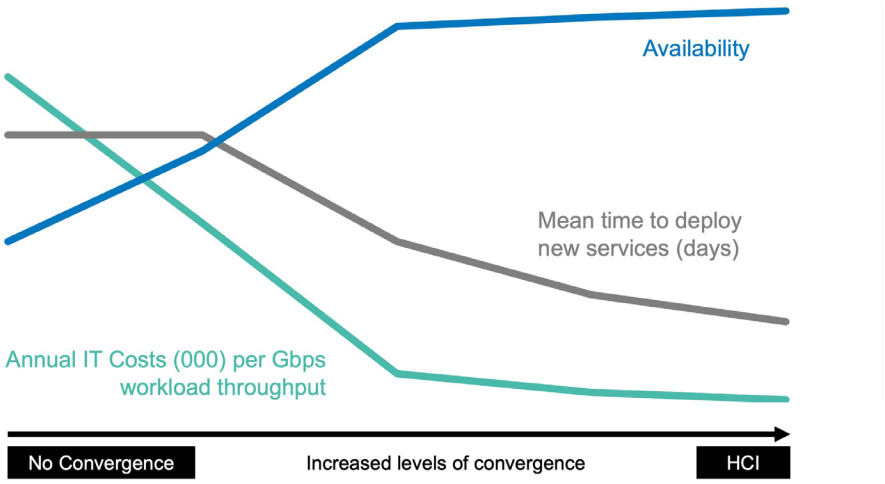

Because HCI is preintegrated, IT organizations no longer need to integrate servers and storage. And it's not necessary to configure hardware and software for specific workloads, because HCI is preconfigured for specific clouds, hypervisors, databases, and virtual desktop infrastructure (VDI) platforms. And for organizations that prefer a cloud experience over an own-and-operate experience, HCI as a service can be implemented with a subscription.

Comparing Traditional and Hyperconverged Infrastructures

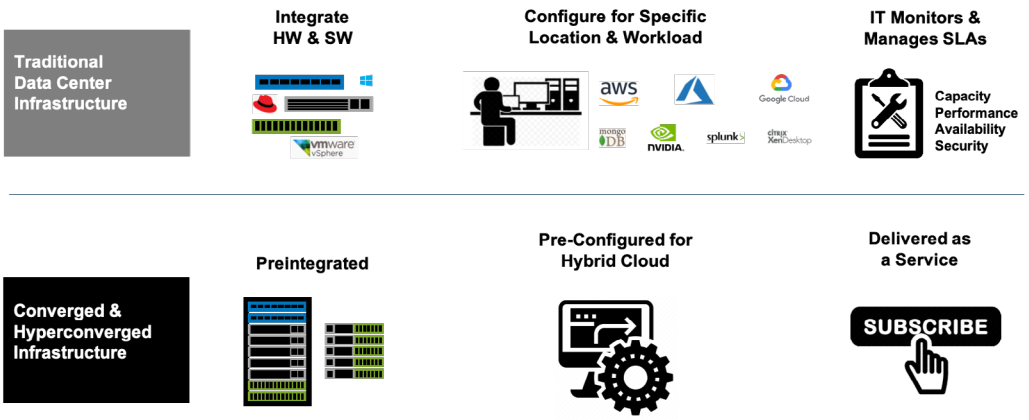

Two HCI Foundations: Aggregated and Disaggregated Systems

HCI promises to free IT pros from the mundane task of implementing data infrastructure so they can spend time on strategic initiatives, like increasing the efficiency of their hybrid cloud and leveraging AI-based analytics.

The next step in fulfilling the promise is understanding the two types of technology behind the solutions: aggregated HCI and disaggregated HCI.

Aggregated HCI

Compute, storage, and networking interfaces are preintegrated by the vendor into an HCI server building block. Operating systems and workload-specific software such as hypervisors, databases, and VDI applications can also be preconfigured. Environments scale by adding server and storage building blocks.

Disaggregated HCI

Disaggregated HCI is like aggregated HCI, except that separate server and storage building blocks are preintegrated and preconfigured by the HCI vendor. This approach increases compute and storage utilization by allowing them to scale independently.

Aggregated and Disaggregated HCI

	Traditional Infrastructure (Nonconverged)	Hyperconverged Infrastructure (HCI)	Disaggregated HCI
Building Blocks	Servers, storage, and networking systems, operating SW, and workload-specific SW	Servers with compute, storage, operating SW, and workload-specific SW integrated inside server	HCI with servers (compute, operating SW, and workload-specific SW) and storage that scales independently
	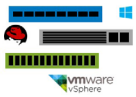		
Integrated by	User	Vendor	Vendor

The Three Big Advantages of Disaggregated HCI: Higher Utilization, Performance, and Capacity

Advantage #1: Higher Utilization of Compute and Storage Resources

The following table shows the number of compute and storage nodes in live NetApp® HCI installations. In no case did they scale the same, and in most cases the difference in the number of compute and storage nodes was at least 3 times. Disaggregated HCI allowed these customers to deploy resources only where they were needed.

In Real Life, Compute and Storage Scale Differently

Actual Use Cases	VDI Private Cloud	Splunk	Kubernetes Private Cloud	VDI Private Cloud	VDI Graphics
HCI Compute	28	62	6	4	2
HCI Storage	8	14	4	6	6

Advantage #2: Higher Performance

HCI servers must support compute and storage overhead. Adding storage steals processing power that was previously available for compute. According to the Gartner Magic Quadrant, the storage nodes for NetApp disaggregated HCI use the industry's leading all-flash array technology. The processing power of these nodes is 10% dedicated to storage, and they are highly optimized for scaling performance with heterogenous and I/O-intensive workloads. Unlike aggregated HCI, disaggregated HCI storage nodes scale performance even with latency-sensitive workloads and when encryption at rest is enabled.

Comparing Aggregated and Disaggregated HCI Performance

Performance Capability	Aggregated HCI	Disaggregated HCI
CPU Overhead	Server processors support overhead of OS, hypervisor, and software-defined storage.	30% less because storage nodes have their own dedicated processors.
Quality of Service	Typically no storage QoS. Noisy neighbor applications disrupt performance of other apps that are accessing data on the HCI server.	Powerful QoS dedicated storage access by fine-grained policies.
Scalability	With large processing overhead, and without storage QoS, scalability is limited.	With less processing overhead and power QoS, scalability is high.

Advantage #3: Higher Effective Capacity

Another benefit of using disaggregated HCI, —with industry-leading array technology, is superior deduplication and compression capabilities. In a side-by-side comparison, the capacity lost to overhead, and the capacity gained through deduplication and compression, were tabulated to find the effective capacity of a 16-node HCI cluster. The effective capacity of the NetApp disaggregated HCI storage nodes was found to be 15% to 73% higher than comparably equipped clusters using traditional HCI technology.

Comparing Aggregated and Disaggregated Effective Capacity

16-Node Solution	Leading Aggregated HCI Vendor A	Leading Aggregated HCI Vendor B	NetApp Disaggregated HCI
Physical space	111.77TB	163.07TB	142.85TB
Effective space after overhead, compression, and deduplication	167.66TB	142.69TB	247.3TB
Percent of physical storage available	150%	87%	173%

Comparing Three System Architectures: Two Based on Aggregated HCI and One on a Disaggregated HCI Foundation

The previous section explained the different foundations for an HCI architecture. This section completes the story by describing the features and benefits of three HCI architectures built on either an aggregated or a disaggregated HCI foundation.

#1: Storage in a VM Architecture

Virtual machine (VM) architecture is found in products from vendors like Nutanix and HPE SimpliVity. It features an aggregated HCI foundation with compute and storage in a single server. With all the hardware preintegrated and with the OS and hypervisor software preconfigured, the systems deploy rapidly. This architecture is distinguished by the fact that business applications and the software-defined storage stack run in virtual machines. However, the scalability of this system performance is inherently limited due to the overhead of the virtualization layer, and because the system processors must support both compute and storage.

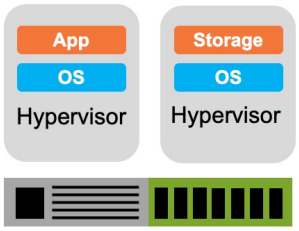

#2: Integrated Hypervisor and Storage Architecture

Products such as VMware vSAN, VMware VxRail, and Microsoft Storage Spaces Direct incorporate the integrated hypervisor and storage architecture. The design features a storage stack integrated into the hypervisor layer. This architecture shares the benefits and limitations of the "storage in a VM" architecture. It delivers fast time to deployment, but performance scalability is limited because of the overhead of virtualization and the need to support both compute and storage.

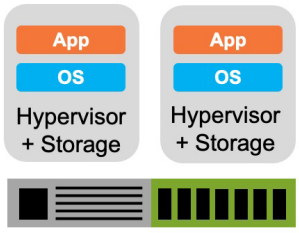

#3: Independent Hypervisor and Storage

This new independent hypervisor and storage architecture is delivered in products from NetApp and HPE Nimble. The design leverages powerful, sophisticated, and mature storage array controller technology for HCI storage nodes. HCI systems that feature this architecture deploy rapidly, with all the server and storage hardware preintegrated and with the server and storage OS and hypervisor software preconfigured. The major difference between this architecture and the aggregated HCI-based architectures is the higher utilization, performance, and capacity that only a dedicated storage controller can deliver.

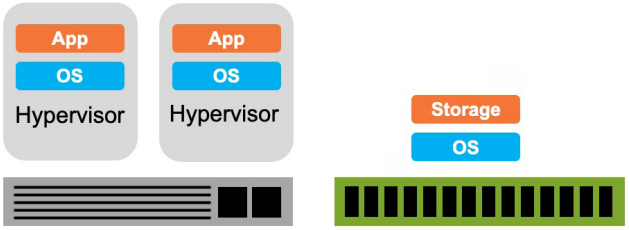

The Bottom Line
Side-by-Side Comparison and Summary

The following table summarizes the capabilities of the three different HCI architectures.

The bottom line is that all HCI systems share the ability to free IT organizations from the mundane tasks of integrating and configuring compute and storage, so that they can invest their precious time in more strategic initiatives. In addition, all aggregated HCI architectures are created somewhat equal, while a disaggregated architecture stands out for its ability to scale compute and storage independently. Disaggregated HCI also stands out as array technology that scales performance higher, delivers more effective capacity, and maximizes the investment in HCI infrastructure.

HCI Architecture Comparison

Architecture	Storage in VM	Storage Integrated in Hypervisor	Independent Hypervisor and Storage
Foundation	Aggregated HCI – Compute and storage share processors	Aggregated HCI – Compute and storage share processors	Disaggregated HCI – Dedicated processors for compute and storage
Time to Deployment	Fast with preintegrated infrastructure and preconfigured software	Fast with preintegrated infrastructure and preconfigured software	Fast with preintegrated infrastructure and preconfigured software
Effective Capacity	Lowest based on empirical testing of deduplication and compression	Mid-range based on empirical testing of deduplication and compression	Highest based on empirical testing of deduplication and compression
Performance Scalability	Limited because HCI server processors are shared by compute and storage	Limited because HCI server processors are shared by compute and storage	Highest because compute and storage processors are dedicated
Compute and Storage Utilization	Lower because compute and storage must scale together	Lower because compute and storage must scale together	Higher because compute and storage can scale independently
Examples	Nutanix and HPE SimpliVity	VMware vSAN, VMware VxRail, and Microsoft Storage Spaces Direct	NetApp HCI and HPE Nimble

About NetApp

NetApp is the leader in cloud data services, empowering global organizations to change their world with data. Together with our partners, we are the only ones who can help you build your unique data fabric. Simplify hybrid multicloud and securely deliver the right data, services, and applications to the right people at the right time. Learn more at www.netapp.com.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

WP-7266-0720

NetApp®