



Not All Hyper Converged
Infrastructure Solutions
Are Created Equal
**A Guide to Understanding
Common HCI Architectures**





Hyper Converged Infrastructure Is a Key Enabler of Digital Transformation	3
Understanding Different HCI Architectures	3
Hypervisor with Storage in a VM	4
Integrated Hypervisor and Storage	5
Independent Hypervisor and Storage	5
Strengths and Weaknesses of Each HCI Architecture	6
Storage in a VM	6
Integrated Hypervisor and Storage	6
Independent Hypervisor and Storage	6
Introducing NetApp HCI	7
Which Solution Will You Choose?	8

Hyper Converged Infrastructure Is a Key Enabler of Digital Transformation

As your enterprise tries to adapt to the digital era, IT infrastructure approaches that worked well just a few years ago may no longer be adequate. To succeed, your IT team needs infrastructure that is easier to manage and that is faster to procure and to deploy. Infrastructure that is built with simplicity in mind frees up your time and budget, enabling greater focus on new applications and services.

Many IT teams are discovering that hyper converged infrastructure (HCI)—which brings together servers, storage, virtualization, and management—is proving to be a better infrastructure choice. HCI can address a broad range of use cases, including end-user computing, private cloud, and consolidation of business-critical databases and applications. Gartner predicts that by 2020, at least 20% of business-critical applications will have moved from traditional infrastructure to HCI.¹

HCI solutions offer significant business benefits that include:

- Ease of purchase
- Simple setup
- Consolidated management
- Full-stack support
- Pay-as-you-go economics

Because of these advantages, HCI enables your IT team to respond more rapidly to new business demands, a crucial advantage for enterprises today.

Although HCI solutions from different vendors offer a similar set of core business benefits, there are some significant architectural differences between solutions. As you evaluate HCI options to meet your next-generation data center needs, it is important to understand the differences between various HCI offerings before you choose a solution. Each architecture has particular strengths and limitations that might make it more or less suitable to satisfy your unique business needs.

This guide examines the characteristics of three HCI architectures and explores the potential advantages and disadvantages of each approach. This information can help you make a more informed decision that closely matches your business needs.

Understanding Different HCI Architectures

HCI solutions combine servers, storage, and virtualization in scale-out building blocks, as illustrated in Figure 1. By the time that HCI solutions were introduced, hypervisor software had already reached a high level of functionality. Therefore, the biggest architectural challenge for HCI implementations was how best to use storage that is distributed across multiple server nodes to create a resilient storage pool that all nodes in a cluster can access. The approach to storage is the biggest architectural differentiator between HCI solutions from different vendors.




Figure 1) HCI replaces traditional IT infrastructure with simpler, scale-out building blocks.

The three approaches to HCI are:

- Hypervisor on bare metal with storage in a virtual machine
- Hypervisor with integrated storage
- Independent hypervisor and storage

Each of these approaches is explained in more detail in the following sections. A later section explores the relative strengths of the different approaches.

¹Gartner Magic Quadrant for Hyperconverged Infrastructure, 2018.

Hypervisor with Storage in a VM

The first approach to HCI (see Figure 2) relies on a special storage VM, sometimes called a controller VM or CVM. With this architecture, each node in the HCI cluster runs a standard hypervisor on bare metal, similar to any other virtualized environment. What separates the HCI solution from a standard virtualized environment is that each hypervisor instance has a dedicated storage controller VM that runs continuously to provide storage services to the cluster. Resiliency is typically achieved by storing two (or more) copies of data, with each copy on a different node. (This approach is typical for all the HCI architectures that are discussed in this guide.)

This architecture can be implemented by using an existing storage operating system (storage OS) inside the controller VM, reducing time to market. Much of the innovation and unique intellectual property in this architecture is found in the installation and management software. The first HCI solutions to reach the market used this approach, and it remains common. The best-known examples are Nutanix and HPE SimpliVity.




Figure 2) The “storage in a VM” HCI architecture relies on a dedicated storage VM on each node to provide storage services across the cluster.

Integrated Hypervisor and Storage

The second approach to an HCI architecture (see Figure 3) integrates storage functions directly into the hypervisor without the use of a CVM. For a vendor to take this approach, the vendor obviously needs direct control of a hypervisor. The only vendor that has taken this approach so far is VMware.

VMware vSAN integrates directly with the VMware ESXi hypervisor. The strength of this approach is in the integration with the latest features of the VMware environment. However, because it was built from scratch, VMware vSAN is still relatively young and has not yet achieved full feature parity with other storage operating systems.




Figure 3) The “integrated hypervisor and storage” architecture adds storage functions to the hypervisor itself.

Independent Hypervisor and Storage

The third HCI architecture uses independent hypervisor and storage nodes, with all software deployed on bare metal (see Figure 4). By physically separating these functions, this approach provides bare-metal performance for storage while confirming that VMs do not compete for resources with storage operations that run on the same node.

This architecture provides greater flexibility when it comes to scaling. If you need more storage capacity or performance, you simply add storage nodes. If you need more compute, you add compute nodes. NetApp® HCI is one example of this approach.




Figure 4) The “independent hypervisor and storage” architecture runs the hypervisor and storage OS on separate nodes.

Strengths and Weaknesses of Each HCI Architecture

All three of the HCI architectures that have been described are easy to buy, simple to install, and offer consolidated management. Beyond those features, the solution that you ultimately choose might depend on your specific business requirements. The Venn diagram in Figure 5 illustrates the relative strengths of each architecture and where those strengths overlap. Table 1, at the end of this section, summarizes these strengths.




Figure 5) Relative strengths of each approach to HCI.

Storage in a VM

The primary advantage of running the storage OS inside a VM is that it enables vendors to get to market quickly while maintaining compatibility with popular hypervisors. This approach also has one or two advantages in common with each competing architecture.

First, it offers the smallest entry footprint and broad platform compatibility, strengths that it shares with the integrated hypervisor and storage approach. If you are targeting remote and edge deployments of fewer than 100 VMs, this architecture provides a low-cost solution in a small footprint.

Similarly, if you need a solution that runs on your preferred hardware platform, you might be better off choosing a solution that uses a storage VM or one that integrates storage with the hypervisor. With those approaches, options are available on a wide range of hardware.

Integrated Hypervisor and Storage

The advantage of an HCI solution that integrates storage with the hypervisor is in the feature integration. For example, VMware owns and controls vSphere and vSAN. Therefore, an HCI solution that is built on that combination of technologies delivers faster and more complete integration with the latest innovations in the VMware ecosystem than competing HCI solutions can.

Another advantage of this approach, one that it shares with an independent hypervisor and storage architecture, is that storage software runs against bare metal. This feature provides direct access to hardware features and delivers greater performance than a storage OS that runs inside a VM can deliver.

Many enterprises today rely on multiple hypervisors, and this approach obviously locks you into a single hypervisor. If you need an HCI architecture that can support different hypervisors, now or in the future, you might prefer to look at alternative HCI solutions.

If your organization is a VMware shop today and you are considering an HCI deployment that is built on vSAN, you should explore the known limitations of the current release of vSAN.

Independent Hypervisor and Storage

By separating nodes that run VMs from nodes that run storage, this architecture is flexible and scalable. You can add compute and storage resources according to your needs, and you have the potential to scale to a higher total node count without diminishing returns. You get predictable performance from the storage and compute nodes that you purchase, without the HCI tax on CPU and on memory that other HCI architectures incur. Because storage and VMs run on separate nodes, the potential for resource contention is greatly reduced, providing more predictable performance.

The separation of storage from compute nodes also provides significant benefits in terms of licensing. For example, when you have an architecture that runs storage in a VM or as part of the hypervisor, a portion of your hypervisor licensing payment goes to pay for running the storage workload. Suppose that each storage VM consumes 20% of the resources on each node. You will end up requiring more nodes to meet your compute needs, and you will pay higher hypervisor licensing costs as a result.

The licensing penalty can also be significant when you consider other software licenses, such as for databases. Popular databases such as Oracle and Microsoft SQL Server are typically licensed per processor. With separate storage nodes, you do not pay licensing fees on the processors that are running storage functions.

By lowering hypervisor and database licensing costs, you reduce your TCO for this architecture.

Introducing NetApp HCI

Like many vendors, NetApp initially considered the “storage in a VM” approach for HCI. However, although it would have been straightforward to encapsulate the NetApp SolidFire® Element® OS—our proven, scale-out storage OS—in a VM, it was clear that we could not deliver its full value that way. Architecturally, it simply made more sense to package Element OS on bare-metal storage nodes. With that approach, you can take full advantage of its all-flash architecture; predictable performance with quality of service; and enterprise storage features, including inline deduplication and compression across the entire cluster.

As part of the enterprise NetApp HCI solution, we leveraged the proven hypervisor-clustering capabilities of VMware ESXi on compute nodes and created a simplified installation routine to get complete HCI systems up and running quickly. We also built a consolidated management UI, enabling you to fully leverage management technologies that you are already using, such as VMware vCenter and vRealize orchestration. Automated deployment of additional hypervisors on NetApp HCI is under development.

Although other HCI vendors have put their emphasis on management, the NetApp HCI architecture focuses on delivering greater flexibility in the compute and storage layers (see Figure 6). NetApp HCI delivers all the benefits that have been described for HCI architectures with independent compute and storage. In particular, NetApp HCI:

- Enables you to more closely match your compute and storage needs
- Offers exceptional scaling
- Reduces hypervisor, database, and other software licensing costs
- Provides predictable performance for greater levels of consolidation




Figure 6) NetApp HCI combines a flexible design with predictable performance, simplified operations, and superior cloud integration.

Selection Criteria	VM	Hypervisor	Independent
Time to market	•		
Lowest entry price	•	•	
Platform compatibility	•	•	
Hypervisor feature integration		•	
Bare-metal performance		•	•
Multiple hypervisors	•		•
Flexibility			•
Scaling			•
Decreased licensing costs			•
Lowest TCO			•
Predictable performance			•

Table 1) HCI architecture capability comparison.

The guaranteed performance of NetApp HCI prevents noisy neighbors and runaway processes from interfering with other applications that run on the same HCI cluster. This benefit makes it an optimal architecture for private cloud and mixed workload environments. Performance is managed automatically, and NetApp gives you the tools to instantaneously address any performance problems that arise. You can allocate capacity and performance independently for every application, and you can easily adjust allocations as workloads shift or as your needs evolve.

An HCI solution must integrate easily with your existing IT operations, both on the premises and in the cloud. Otherwise, it becomes another infrastructure silo, making your data center more complex. In a next-generation data center, you must be able to manage and to protect data globally. You also must integrate with other important applications and services in your data center environment and beyond.

By delivering predictable performance and simplified operations on a highly flexible and efficient cloud architecture, NetApp HCI increases the agility of your business. NetApp HCI is Data Fabric ready out of the box, so you can access all your data across any cloud: public, private, or hybrid. Because data is accessible both on the premises and in the cloud, the Data Fabric integration in NetApp HCI enables your company to respond and innovate more quickly.

Which HCI Solution Will You Choose?

All the HCI architectures that are discussed in this guide satisfy the basic expectations of HCI by simplifying purchasing, deployment, and management. Now that you understand the differences between these architectures, it should be clear that each approach has both strengths and weaknesses. Use the Venn diagram in Figure 5 and the list of selection criteria in Table 1 to help you zero in on the best HCI solution to satisfy your specific requirements.

For a detailed technical discussion about the three architectures, watch a video of a recent Tech Field Day presentation, [NetApp Comparing HCI Architectures](#). To learn more about NetApp HCI, visit the [NetApp HCI page](#) on [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.

Copyright Information

Copyright © 1994–2018 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

WP-7266-0418