

How Architecture Design Can Lower Hyperconverged Infrastructure (HCI) Total Cost of Ownership (TCO)

By Eric Slack, Sr. Analyst

December 2017



Evaluator Group

Enabling you to make the best technology decisions



Evaluator Group



NetApp®

Total cost of ownership (TCO) in IT systems is a function of their capacity to do work, which in turn, is a function of storage and compute performance. In a Hyperconverged Infrastructure (HCI), this capacity is expressed in terms of the number of virtual machines (VMs) that can be supported for any given configuration. TCO is also affected by architectural design, as HCIs combine storage and compute resources into a scale-out cluster.

In this TCO analysis we calculate the storage and compute performance of NetApp HCI and two leading competitive systems and compare how their different architectures impact total cost at varying levels of scale.

Benefits of Hyperconverged Infrastructure

HCI technology combines compute power with local storage and storage networking, running one of the available hypervisor platforms in a clustered architecture. A turnkey solution, HCIs have been popular for specific applications that required a new or dedicated infrastructure such as VDI or for isolated environments, like branch offices or remote departments.

As a scale-out solution, HCIs are designed to be deployed quickly, be simple to operate and be easily expanded, making them a good fit for IT generalists or even non-IT personnel, in some cases. These comprehensive infrastructures are flexible enough to support a variety of different applications and can be configured to meet the specific resource requirements of a diverse workload mix – to a certain extent.

Challenges of Hyperconverged Infrastructure

Hyperconverged clusters expand easily, but the mean unit of scale is the node itself. And, often nodes in the cluster must have the same storage configurations. This lack of resource flexibility runs counter to a primary benefit of the hyperconverged infrastructure, the ability to support a wide variety of workloads.

Companies want to consolidate applications, and the infrastructure supporting those applications, into an HCI cluster (one of the findings of a study conducted by Evaluator Group, “HCI in the Enterprise”). This consolidation creates a non-homogeneous collection of workloads that can require a dynamic mix of compute and storage resources. Unfortunately, most HCI clusters force you to add these resources in lock-step with each other, based upon the configuration of the appropriate nodes. But in addition to scaling, there’s a more fundamental challenge in sizing HCI that creates inefficiency and increases cost.

Storage and Compute Performance

HCI have two performance components, storage and compute, that must be supplied in sufficient amounts to support a given number of application VMs. Storage performance is a measure of a system's ability to supply storage IO to a VM workload within an acceptable response time, based on the storage devices and storage software stack. Compute performance is the system's ability to process program instructions based on the product of a system's clock speed and available CPU cores, for a given class of processors. For more information on the testing process, see the IOmark link in the next section.

Compute and storage performance should be calculated independently for HCI systems since they represent different performance capabilities and are driven by different subsystems within the hyper converged architecture. The most appropriate metric for expressing this "performance capacity" is the number of VMs supported by each node, since it allows both storage and compute performance requirements to be represented in a single per-node spec. The following section explains how storage and compute performance were determined for this TCO comparison.

Explanation of Performance Testing

In order to evaluate storage performance the **IOmark** benchmark tool was used, with the IOmark-VM workload, which calculates the number of VMs supported for a given configuration. IOmark-VM uses real-world workloads based on the **VMmark** spec developed by VMware. These server virtualization workloads are a composite of application servers typically found in a modern data center, including the following:

- Exchange database server with a 1,000-user workload running on the Exchange server.
- DVD store database application with three DVD store web application servers generating workloads
- Olio database and Olio web application servers
- Hypervisor automation workloads

For storage performance, the two Competitor HCI systems were found to support 80 VMs per node and the NetApp HCI 230 VMs per storage node, as shown in Table 1. This almost 3:1 differential is one factor for the NetApp HCI's lower total cost. (see section below "Why NetApp HCI Costs Less").

Compute performance was determined thorough an analysis of published VMmark tests for the server systems employed by the three HCI products compared (see Table 1). The NetApp HCI Compute nodes, which use an 18-core processor, were found to have lower compute performance than the two Competitor systems which use 22-core CPUs. The specs for these systems are listed below.

HCI System	Storage Performance	Compute Performance
NetApp HCI	230 VMs per node	88 VMs per node
Competitor 1	80 VMs per node	128 VMs per node
Competitor 2	80 VMs per node	128 VMs per node

Table 1

HCI Systems Used in this Comparison

Both of the Competitor HCI systems were comprised of a cluster of 2U nodes. Each node contained dual Intel E5-2699 v4 CPUs with 22-cores, 512GB of memory and 3 TB of all-flash storage. The cluster was scaled by adding single nodes, from a minimum of four nodes.

NetApp HCI consisted of a cluster of Compute nodes and Storage nodes, each occupying one-half of 1U rack space. The Compute nodes were comprised of dual Intel E5-2695 v4, 18-core CPUs with 768 GB of memory. The Storage nodes were comprised of 8-core CPUs with 384 GB of memory and 2.88 TB of flash storage. The cluster was scaled by adding single nodes from a minimum of four Storage and two Compute nodes.

For this comparison we used VMware vSphere ESXi 6.0 on all three HCI systems.

Data reduction, compression and deduplication, increase effective storage capacity but can also reduce storage performance. Both Competitor systems had compression and deduplication disabled for the storage tests. Conversely, in the NetApp HCI in-line compression, in-line deduplication and post-process compression are always enabled. To provide a direct comparison, both competitors would need to enable compression and deduplication, most likely further reducing the number of VMs per node supported.

For an 800-VM environment the configurations shown in Figure 1 are required, based on the performance numbers from Table 1.

Cluster Configurations for 800-VM Environment




Figure 1

Both the Competitor system clusters have a total of ten nodes, each node occupying a 2U chassis. The NetApp cluster is comprised of fourteen nodes, each considerably smaller, with four nodes in each 2U chassis. The bottom chassis has the bezel removed to show two Storage and two Compute nodes, as an example.

For the total deployed HCI cluster, the data center footprint was reduced from 20U to 8U with NetApp HCI, a reduction of 60%. Acquisition cost is also lower (see below) for NetApp HCI. Although the NetApp cluster has 14 nodes, both Storage and Compute nodes are less expensive than the combination HCI nodes of the Competitor systems, resulting in a 25% lower acquisition cost than Competitor 1 and a 57% lower acquisition cost than Competitor 2.

The total cost of ownership (TCO), shown in Table 1 below, for each of the HCI solutions tested is comprised of the following components:

Acquisition - includes the manufactures' list price of each HCI node, one VMware vSphere Enterprise license for each node, any management software licenses not included in the base node price, deployment costs plus cables.

Support - comprised of applicable hardware and software maintenance costs for one year.

Overhead - includes Administration, Power and Space:

Administration is comprised of cluster management and hypervisor management. Cluster management requires one hour per week for admin support of the hardware infrastructure (including storage), regardless of the number of nodes deployed. Hypervisor management requires one-half hour per week for each node that runs a hypervisor. This is every HCI node for Competitor HCIs and just the Compute nodes for the NetApp HCI. A cost of \$75 per hour was used for an IT admin (based on a fully burdened cost of \$150,000 per year).

Power was calculated at \$0.15 per kWh, based on 300W per NetApp nodes (Storage or Compute) and 450W per HCI node.

Space was calculated at \$2000 per 42U rack per month.

Competitor #1

Cluster of 10 HCI nodes as described above

Acquisition	\$879,830
Support	78,033
Overhead	40,661
Total – 800 VMs	\$998,524

Competitor #2

Cluster of 10 HCI nodes as described above

Acquisition	\$1,526,160
Support	62,848
Overhead	40,661

Total – 800 VMs	\$1,629,669

NetApp HCI

Cluster of 4 Storage nodes and 10 Compute nodes

Acquisition	\$ 655,240
Support	71,599
Overhead	32,843

Total – 800 VMs	\$ 759,682

The TCO for these three systems are shown in Figure 2. It should be noted that at list prices, Acquisition (mostly hardware cost) dominates these cost calculations. If a discount is applied, the Support and Overhead components would make up a larger percentage of the total cost.




Figure 2

NetApp HCI is 24% less expensive for an 800-VM environment compared with Competitor 1 and 53% less than Competitor 2. When the infrastructure is scaled to support more VMs, this cost differential grows, as shown in Figure 3.




Figure 3

When the environment grows the infrastructure has to scale to support more VMs. This growth is based on the storage and compute performance shown in Table 1. NetApp HCI adds a Compute node for every 88 VMs but only needs a new Storage node for every 230 VMs while the Competitor systems must add a new HCI node every 80 VMs.

These per-node performance differences create more cost disparity between HCI solutions. At 1500 VMs the NetApp HCI is 28% less than Competitor 1 and 56% less than Competitor 2. At 2500 VMs the NetApp HCI can save 33% over Competitor 1 and 59% compared with Competitor 2.

Analysis - Why NetApp HCI Has a Lower TCO

Storage Performance

NetApp HCI testing shows their system has 3x the storage performance of the other two systems, on a per-node basis. While the reason behind these performance results is beyond the scope of this TCO analysis *, the manufacturer points out that the SolidFire Element OS software that runs the storage nodes within NetApp HCI was designed to maximize flash performance, unlike other software defined storage technologies that were mostly designed for disk storage.

Better storage performance means the NetApp HCI cluster requires fewer storage nodes to support a given number of VMs. This, in turn, drives lower costs of acquisition and support (software costs are based on the hypervisor, which doesn't run on Storage nodes).

Disaggregated Architecture

The other factor behind its lower total cost is the NetApp HCI architecture. Since most hyperconverged infrastructures combine storage and compute functions in each node, the lower of these two performance numbers becomes the determinant of that product's ability to support VMs. In this case each of the competitor HCI systems is limited to 80 VMs per node based on their storage performance.

NetApp HCI has separate storage and compute nodes so clusters can be built based on the number of *each node type* required, instead of in a fixed ratio of storage and compute. This disaggregated architecture allows a NetApp HCI cluster to take advantage of its storage performance (230 VMs per Storage node) without stranding these resources by scaling based on the lower compute performance. It also allows the cluster to expand storage or compute capacity in a more granular fashion, and better meet the requirements of the aggregate workloads it supports.

Summary / Conclusions

Historically, HCIs have traded simplicity for efficiency. By combining storage and compute resources into the same physical nodes they make deployment and expansion easier, but make sizing and scaling less precise and more expensive. The NetApp HCI architecture separates these two resources so they can scale independently, allowing clusters to be configured with less waste and enabling them to leverage their much greater storage performance. The result is an HCI solution with a TCO between 24% and 59% less than that of two leading competitors.

About NetApp

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit www.netapp.com. #DataDriven

About Evaluator Group

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

Copyright 2017 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.