

Technical Validation

Assuring Database Performance and Availability with NetApp HCI

Scalable, On-premises, Hyperconverged Infrastructure for Microsoft SQL Server

By Tony Palmer, Senior Validation Analyst

November, 2020

This ESG Technical Validation was commissioned by NetApp and is distributed under license from ESG.



Contents

- Introduction 3
 - Background 3
 - NetApp HCI 4
- ESG Technical Validation 4
 - Scalability 5
 - ESG Testing 5
 - Availability and Data Locality Impact 7
 - Guaranteeing Performance 8
 - ESG Testing 8
- The Bigger Truth 12

ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team’s expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

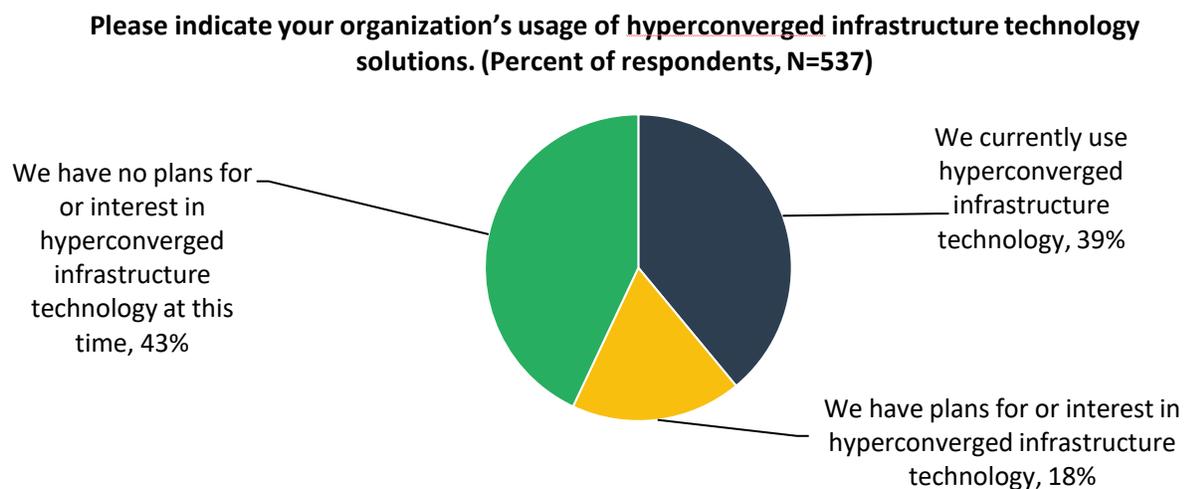
Introduction

This ESG Technical Validation highlights the capability of NetApp HCI to maintain a guaranteed level of database performance while independently scaling compute and storage with no effect on operations. Management simplicity across applications, availability during maintenance, node failure and recovery, and the resource and cost efficiency of NetApp HCI are also examined.

Background

Organizations today must be extremely agile and flexible to add applications to production environments quickly to respond to the needs of the business. This is one reason for the popularity of hyperconverged infrastructures (HCIs) (see Figure 1). HCI offers a single, centrally managed solution with software-defined compute, network, and storage. The performance gains introduced to HCI platforms using all-flash and NVMe technology have opened the door to mission-critical workloads. Today, many organizations successfully run latency-sensitive workloads like Microsoft SQL Server on their HCI clusters. Planning, sizing, provisioning, and configuring a database environment is a complex task. HCI can potentially address all these issues. In an ESG research study, HCI deployment drivers most cited by respondents include improved scalability (31%), total cost of ownership (28%), ease of deployment (26%), and simplified systems management (24%).¹

Figure 1. Hyperconverged Infrastructure Usage Trends



Source: Enterprise Strategy Group

Organizations are told to digitally transform, to become more agile, and to respond to the business faster to survive in a highly competitive market. One method used by organizations to digitally transform is modernizing their infrastructures, which means shifting from a traditional three-tier architecture to a solution that integrates compute, storage, networking, and virtualization. Such a solution must deliver a more cloud-like experience on-premises, making the eventual transition to hybrid multi-cloud easier while enabling organizations to confidently move cloud-native applications between the public cloud and an on-premises private cloud.

In ESG's annual technology spending intentions survey, 24% of organizations said that deploying hyperconverged infrastructure was one of the areas of data center modernization in which they would make the most significant investments in 2020, and 93% of organizations indicated that spending on hyperconverged and converged technology

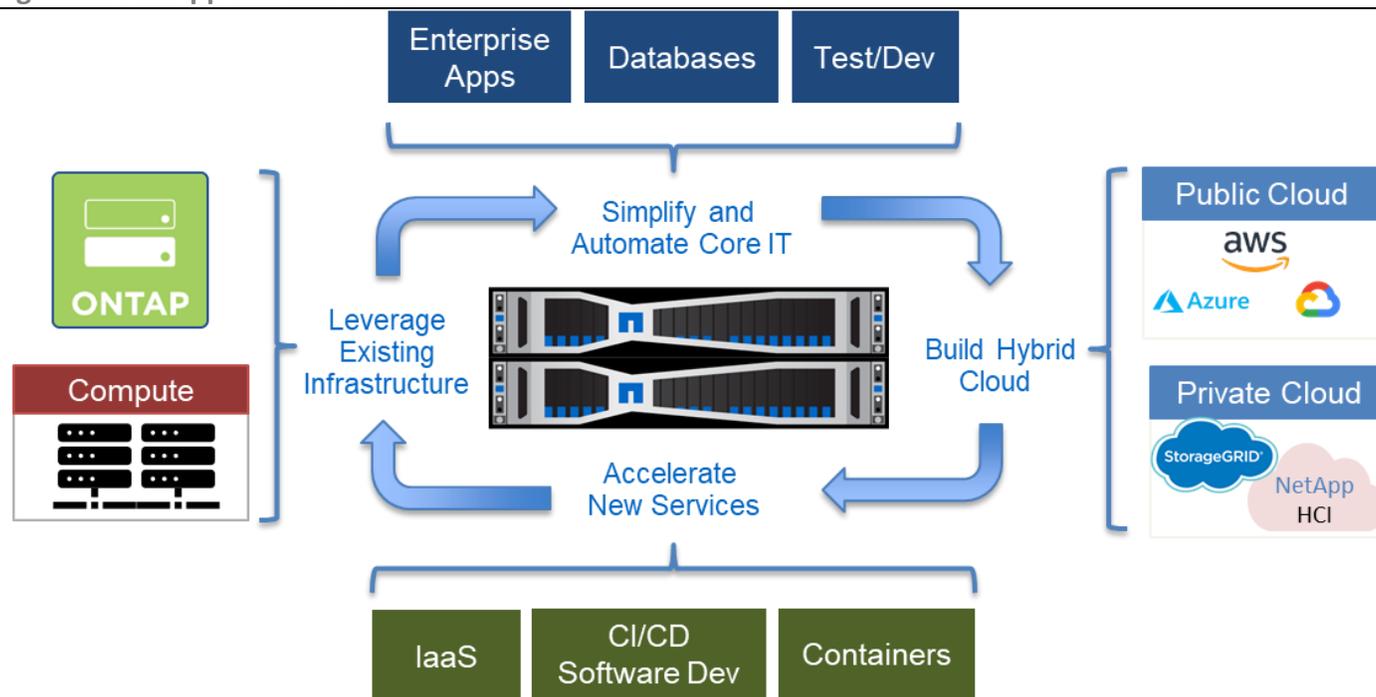
¹ Source: ESG Master Survey Results, [Converged and Hyperconverged Infrastructure Trends](#), October 2017.

would be sustained at current levels or increase in 2020.² This demonstrates that organizations view HCI as an integral part of their modernization strategy, and it is now powering mission-critical production applications.

NetApp HCI

NetApp HCI is a hyperconverged system designed to deliver a public cloud consumption experience with simple management and configuration, independent compute and storage resource allocation for dynamic scale, and predictable performance for mixed-workload hybrid clouds. Enterprise customers use NetApp HCI performance guarantees to consolidate all their applications, including the ones that previously required dedicated silos. However, this is only one of many use cases for NetApp HCI. NetApp HCI is built to seamlessly leverage existing infrastructure, simplify and automate core IT, build out a hybrid cloud environment, and accelerate the development and delivery of new services.

Figure 2. NetApp HCI Use Cases



Source: Enterprise Strategy Group

Powered by NetApp Element software, NetApp HCI scales resources independently to deliver consistent performance. It supports edge-to-edge workloads from mission-critical core data center applications to remote locations. With Element, each volume, datastore, and virtual machine can be configured with minimum, maximum, and burst IOPS values. The minimum IOPS setting guarantees performance, independent of other applications on the system. The maximum and burst values control allocation, enabling the system to deliver consistent, appropriate performance to all workloads. These workload settings are built into the workflow of provisioning new storage or can be modified on the fly via the GUI, a VMware plug-in, or APIs.

ESG Technical Validation

ESG performed evaluation and testing of NetApp HCI in a Microsoft SQL Server environment. Testing was designed to compare SQL Server performance and scalability on NetApp HCI to a conventional HCI platform we called Vendor Z. Independent scaling of storage and compute with no effect on running applications was also tested.

² Source: ESG Research Report, [2020 Technology Spending Intentions Survey](#), February 2020.

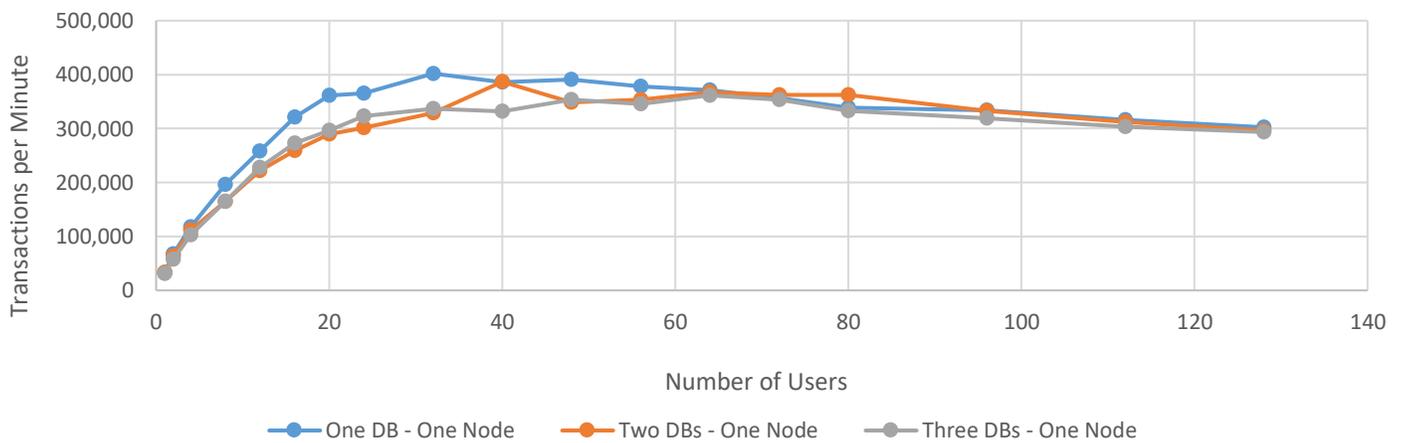
Scalability

Microsoft SQL Server 2017 was used for testing and was installed on a four-node conventional HCI cluster, and a NetApp HCI cluster with four compute and storage nodes. Both clusters used a total 24x 480GB SSDs—six per node. Workloads were identically configured for both environments. HammerDB was used to generate a write-heavy workload consisting of a mix of five concurrent transactions of different types and complexity either executed online or queued for deferred execution against multiple types of tables with a wide range of record and population sizes. The tests simulated up to 128 users on each platform.

ESG Testing

First, we scaled databases and users within a single node, as shown in Figure 3 and Figure 4.

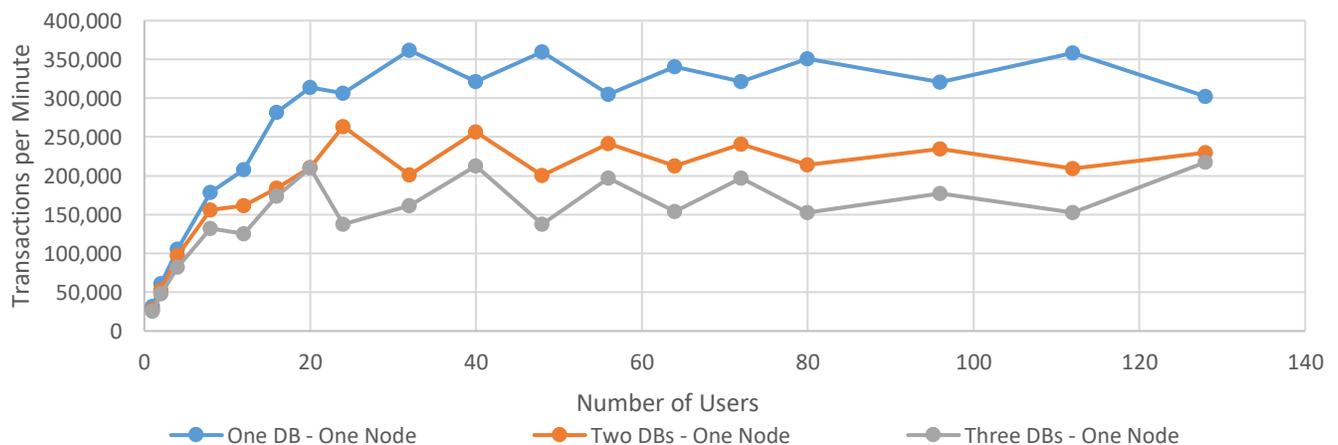
Figure 3. NetApp HCI—Scaling Databases in a Single Node



Source: Enterprise Strategy Group

NetApp HCI was able to sustain reliable, predictable performance as we scaled users, with one, two, or three concurrent databases running on the cluster with minimal variance in performance throughout the test. The conventional HCI system results were illuminating. Performance scaling as users were added was much more variable, and as databases were added to the node, performance decreased significantly by up to 62%

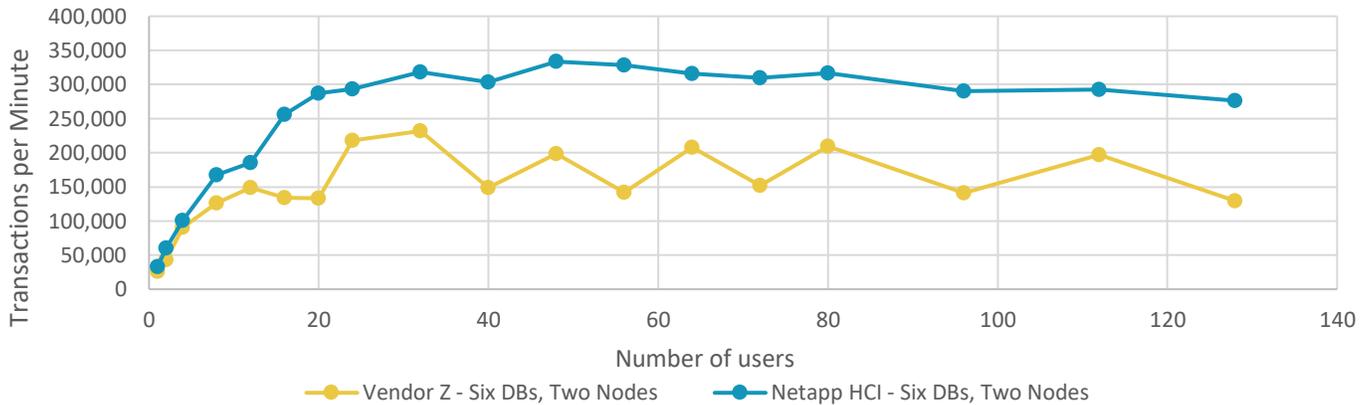
Figure 4. Conventional HCI—Scaling Databases in a Single Node



Source: Enterprise Strategy Group

Next, we scaled up to six databases across two nodes. Results achieved with the two platforms are compared in Figure 5.

Figure 5. Scaling Databases Across Multiple Nodes

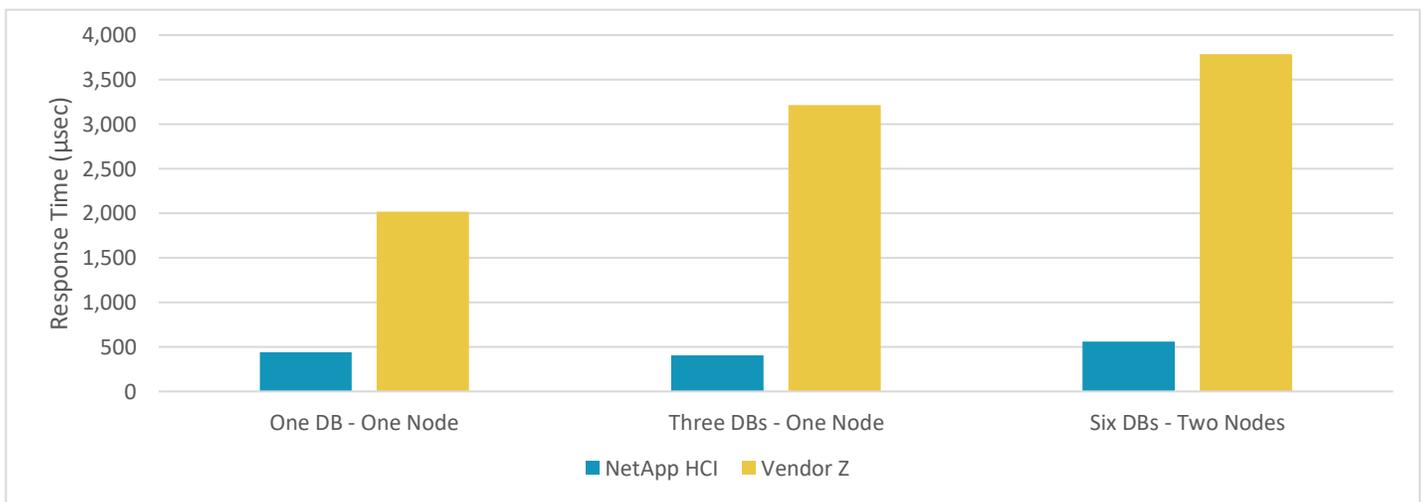


Source: Enterprise Strategy Group

Again, NetApp HCI scaled performance smoothly and predictably as users were added, peaking at 333,517 TPM while conventional HCI performance was again quite variable, and the platform maxed out at 231,888 TPM.

Response time is an extremely important metric when talking about transactional database applications because response time directly impacts application responsiveness and the user experience. Figure 6 shows the average response times for both platforms when scaling from one DB on one node to six DBs on two nodes.

Figure 6. Average Total Response Time



Source: Enterprise Strategy Group

NetApp HCI delivered higher transaction rates in every test we performed with sub-millisecond response times. The conventional HCI platform’s response times across every test were five to seven times higher, on average.

i Why This Matters

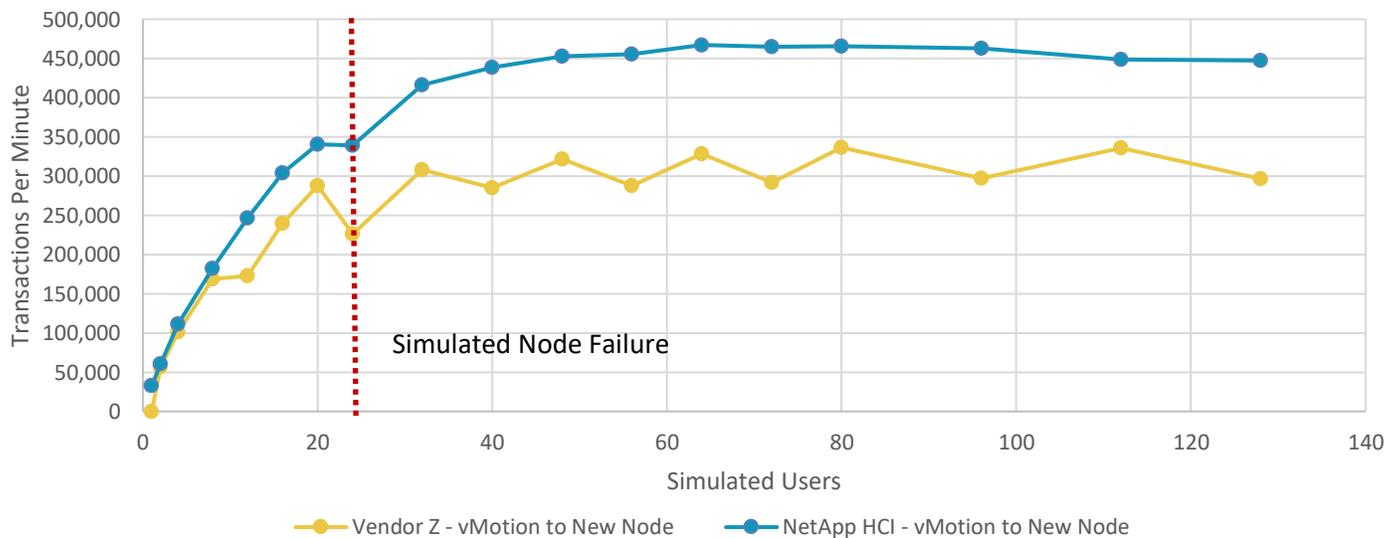
Delivering high levels of performance is a requirement for IT organizations that rely heavily on mission- and business-critical databases. This is especially important in environments where data growth is constant and continuous data access is a requirement. The ability to easily meet these performance and scalability requirements is essential for anyone evaluating hyperconverged infrastructures.

ESG testing validated that NetApp HCI outperformed a similarly configured conventional HCI platform in every metric we examined. NetApp HCI sustained higher transactions per minute, with less variability, at significantly lower response times. This meant reliable, predictable performance as we scaled users and databases running on the cluster with minimal variability in performance throughout the test, which is meaningful in the real world. This can translate into support for a significantly higher density and performance of supported databases and applications.

Availability and Data Locality Impact

NetApp HCI disaggregated architecture enables the platform to separate compute and storage within the cluster. This is different from conventional HCI, where storage and compute are combined on all nodes. ESG looked at the ability of both platforms to sustain performance through a node failure and the impact of data locality. We ran the same SQL Server workload using HammerDB twice on each platform, first with no simulated failure, then in the second run we vMotioned the SQL Server to a new, empty node at the 24-users mark to show the impact of data locality.

Figure 7. Impact of Data Locality During a Node Failure



Source: Enterprise Strategy Group

As seen in Figure 7, conventional HCI suffered a significant impact when the SQL Server was vMotioned to the empty node, processing 26% less transactions than in the test run with no failure. Transactions rose slowly, getting within 6% of the uninterrupted test after about 40 minutes, but never quite reached the levels of the uninterrupted test, averaging -7.5% for the remainder of the run.

NetApp HCI transactions dropped by just 10% when the SQL Server was vMotioned, but performance stabilized after less than 10 minutes, and was within .4% of the uninterrupted test results for the remainder of the run.

Guaranteeing Performance

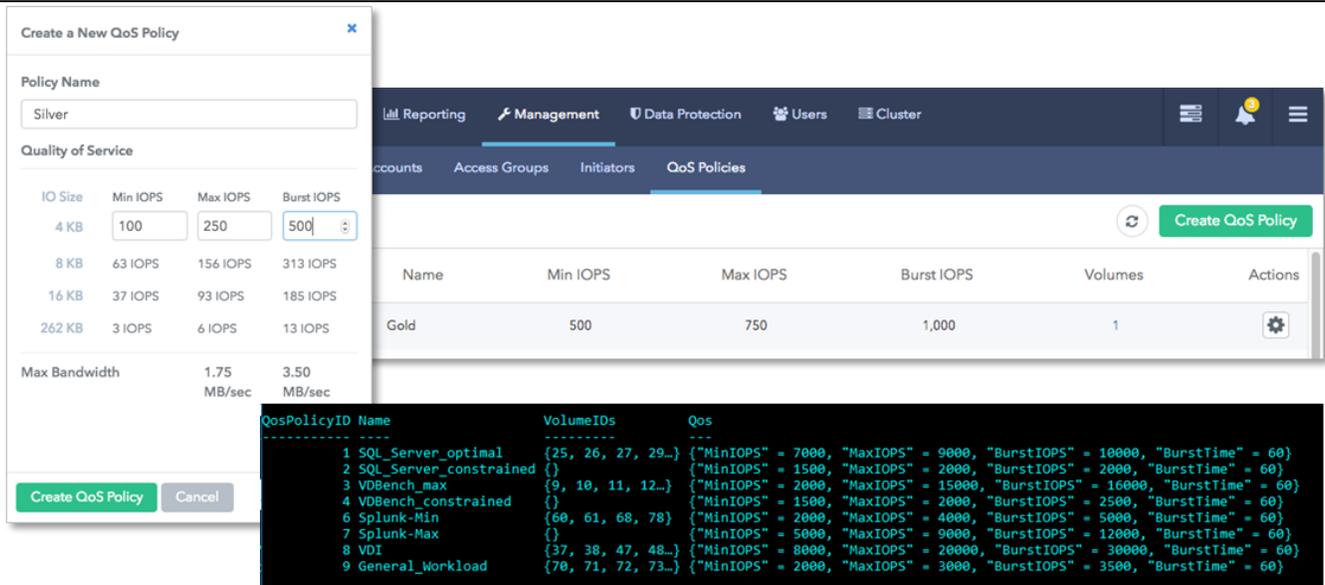
In the virtualized IT world, shared infrastructure can experience a performance imbalance from workloads sharing resources. Critical applications can be starved of storage I/O by another busy workload, resulting in unacceptable response time. NetApp HCI was designed both to limit aggressive, resource-hungry workloads and to guarantee a minimum level of performance to critical workloads. A key benefit is that IT no longer needs to execute tasks such as overprovisioning hardware or splitting workloads among nodes or clusters to ensure adequate performance under any circumstance. All those tasks can increase both complexity and cost.

ESG Testing

ESG tested the ability of NetApp HCI to limit the performance resources that an ill-behaved application consumes while guaranteeing minimum performance for critical apps and compared that to the conventional HCI platform. This multifaceted approach protects critical app performance whether an unruly application is limited or not. The ability to set minimum IOPS for applications is why NetApp HCI can maintain SLAs.

NetApp HCI QoS specifies the performance settings applied to volumes—minimum, maximum, and burst IOPS. When a volume is created with a policy, it is linked to that policy, and updates to the policy are applied to all linked volumes. We assigned all volumes associated with each workload to a different policy. When we needed to modify QoS for an application, we changed the policy and it was instantly applied to all volumes linked to the policy. Policies can be created, assigned, and adjusted by using the NetApp API as well.

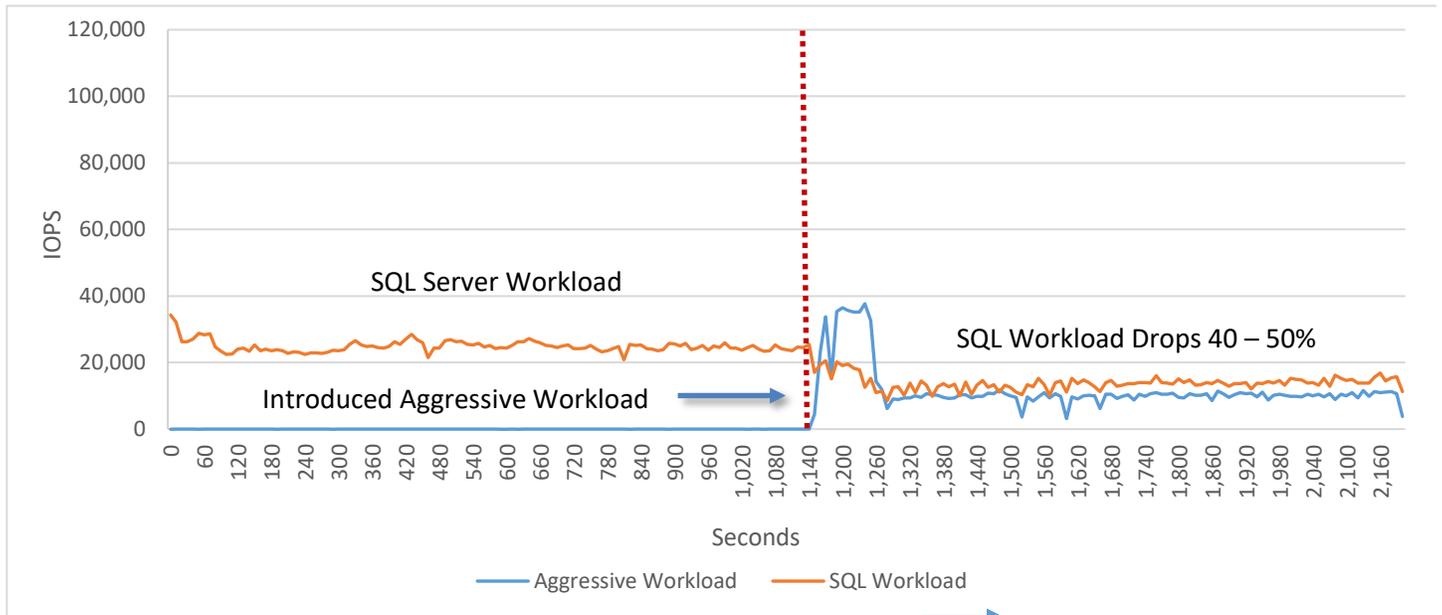
Figure 8. QoS Policies



Source: Enterprise Strategy Group

We ran one SQL Server workload on each node along with an aggressive workload on both platforms. On NetApp HCI, the SQL workloads were guaranteed a minimum level of performance using their linked QoS policy (MinIOPS). Figure 9 shows the results of the conventional HCI test run. The chart shows the aggressive workload as a moving average of 10 data points to make it easier to follow. The SQL workload was processing about 25,000 IOPS with a 5.5ms response time. When we introduced the aggressive workload, IOPS immediately felt the impact and dropped by 50% and stayed between 40% and 50% lower while response time nearly tripled to 14ms.

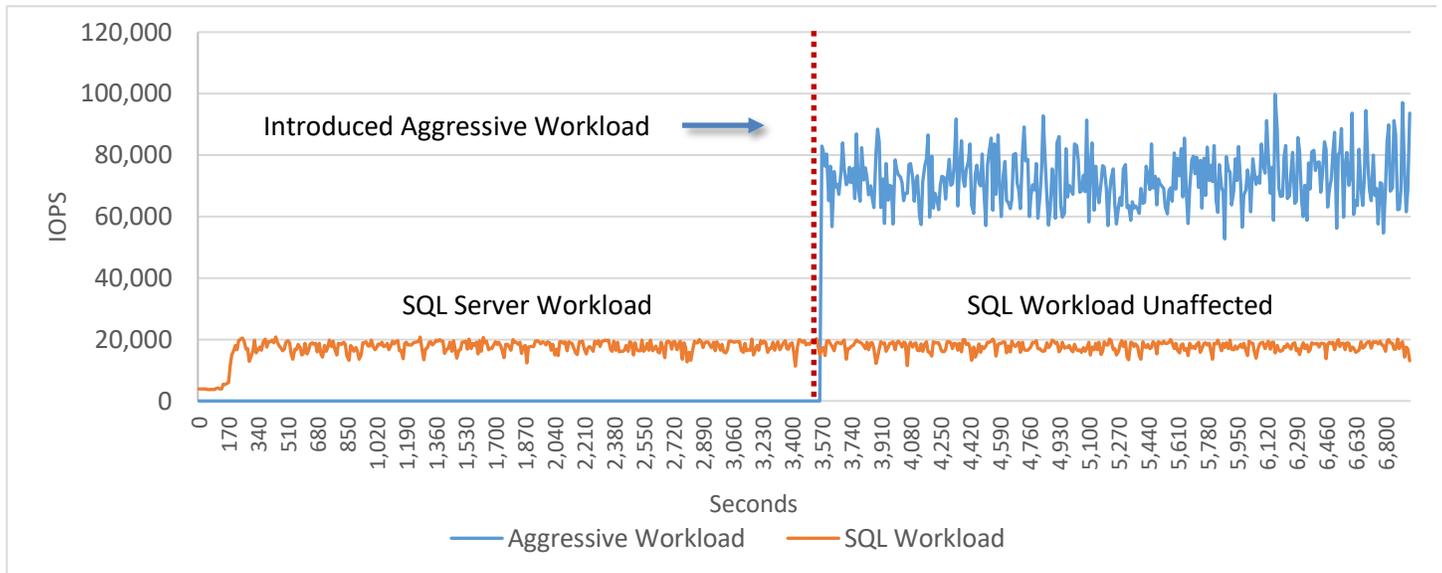
Figure 9. Conventional HCI SQL Server IOPS with Aggressive Workload



Source: Enterprise Strategy Group

Next, we executed the same test run against the NetApp HCI cluster.

Figure 10. NetApp HCI SQL Server IOPS with Aggressive Workload



Source: Enterprise Strategy Group

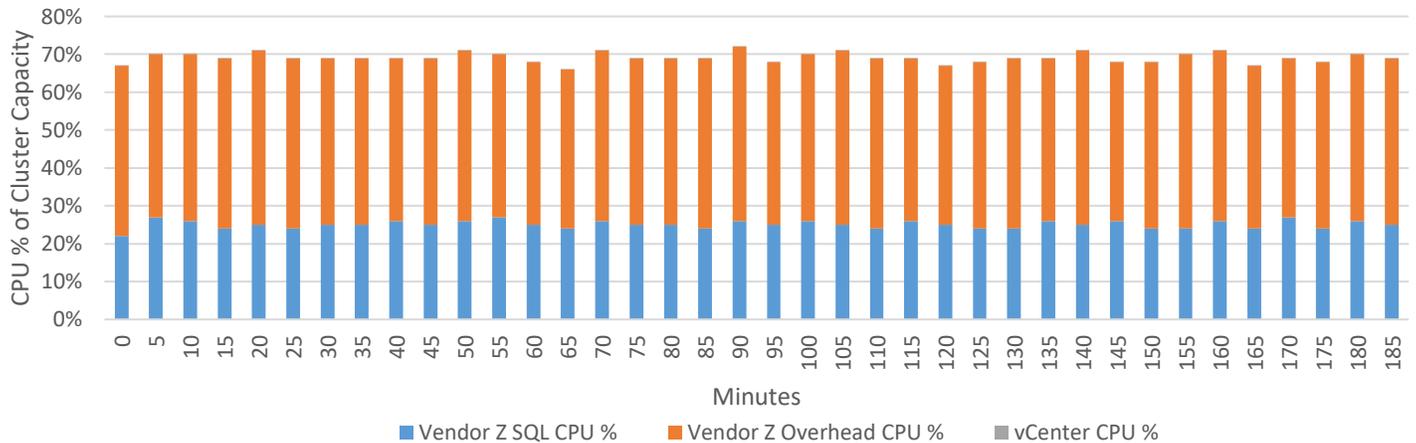
As Figure 10 shows, The SQL workload was driving about 20,000 IOPS with a 2.6 ms response time. When we introduced the aggressive workload, IOPS were completely unaffected. While the response time did increase, the impact was much lower than with conventional HCI. At this point, the unconstrained aggressive workload was consuming all the resources it could, but the SQL Server instances were still completely unaffected thanks to their minimum performance guarantee.

Finally, we compared CPU utilization between NetApp HCI and the conventional HCI cluster. The results were interesting to say the least. VMware reports CPU usage in megahertz (MHz). To compare the utilization across both platforms, we calculated the total MHz available on each cluster by calculating the MHz available on each ESXi host, then aggregating the

MHz of the four hosts in the cluster. We ran the same SQL workload as in the previous tests, with one SQL Server running per node. We monitored CPU utilization using VMware and native tools on the clusters.

Figure 11 shows the CPU utilization of the conventional HCI cluster as a percentage of the total CPU capacity of the cluster. The conventional HCI cluster stayed near 70% utilization throughout the test. It's important to note that overhead—the CPU utilized for internal cluster activities—averaged 44% of the total cluster capacity throughout the test.

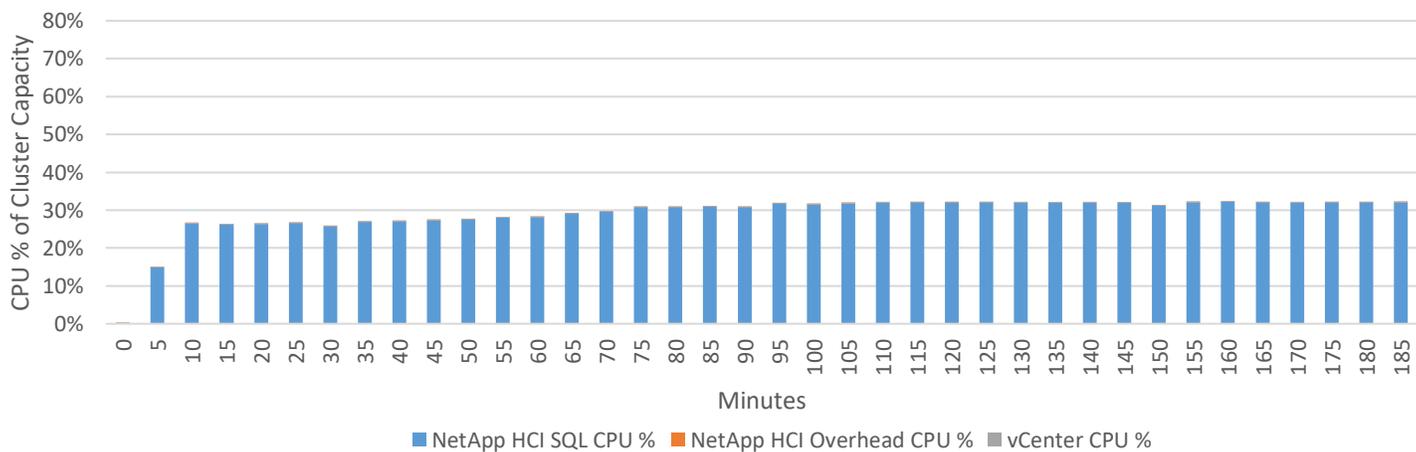
Figure 11. Conventional HCI CPU Consumption



Source: Enterprise Strategy Group

Figure 12 shows the CPU utilization of the NetApp HCI cluster as a percentage of the total CPU capacity of the cluster. We kept the scale of the chart the same as Figure 11.

Figure 12. NetApp HCI CPU Consumption



Source: Enterprise Strategy Group

The NetApp HCI cluster stayed near 30% utilization throughout the test. It's important to note that overhead—the CPU utilized for internal cluster activities—averaged just .05% of the total cluster capacity throughout the test.

Why This Matters

Delivering performance is a key concern for IT administrators when consolidating workloads. Unless they can guarantee performance for high-priority workloads, IT administrators have little choice but to silo databases and applications. That leaves organizations that cannot afford performance degradation during resource contention unable to take advantage of HCI's efficiency. Providing all file and object services in the same node presents a challenge with conventional HCI—there is significant overhead between the essential underlying services that must always be running and the applications themselves. This makes it difficult at best to ensure proper functionality of the hyperconverged infrastructure and meet strict application performance SLAs. Sustaining performance when workloads move around the cluster is also problematic, and an HCI platform should rebalance workloads quickly and efficiently.

ESG has validated that NetApp HCI sustained performance through a simulated node failure with minimal impact, made it simple to prioritize applications with policy-based performance, and demonstrated significantly lower CPU utilization than the conventional HCI platform we tested. Workloads are limited from both overconsuming resources and guaranteed minimum performance levels that can be adjusted on demand. As business needs change and workloads are added to the system, administrators can easily add resources that deliver additional performance for the entire storage pool.

This provides organizations with the cost, scalability, management, and efficiency benefits of HCI, mixing databases and applications of varying priorities, while maintaining predictable performance. They can shift resources in real time and manage scheduled policy changes to support times of high activity using the GUI, CLI, or API. This saves IT time and effort, enabling staff to focus on more strategic activities.

The Bigger Truth

Early HCI solutions were designed to maximize simplicity of deployment and management. But they often lacked the scaling flexibility and performance levels required by today's mission-critical workloads, and organizations risked failure to meet SLAs by combining multiple workloads on HCI nodes. Even so, HCI deployments are gaining in popularity: According to ESG research, 57% of respondents reported that they use or plan to use HCI solutions. This is not surprising because HCI offers numerous tangible benefits: By consolidating infrastructure into software-defined, centrally managed modules instead of compute, network, and storage silos, organizations gain efficiency in both footprint and management, reducing costs and complexity. HCI deployment drivers most cited by respondents include improved scalability (31%), total cost of ownership (28%), ease of deployment (26%), and simplified systems management (24%).³

ESG validated that NetApp HCI provides multiple benefits:

- NetApp HCI delivers performance at scale for SQL Server across diverse use cases with predictable performance and SLA guarantees.
- NetApp HCI seamlessly maintains SQL Server performance and availability during maintenance, node failure, and recovery.
- NetApp HCI eliminates the “HCI Tax” and delivers efficient resource utilization at scale, with minimal CPU overhead, providing financial differentiation and TCO improvements for SQL Server 2017.
- In contrast with conventional HCI infrastructure, which provides file and object storage as proprietary internal components of the cluster—consuming cluster resources and hampering efficiency—NetApp HCI's open Architecture supports best in class file and object storage to enable organizations to minimize resource and cost overhead.

The results presented in this document are based on testing in a controlled environment. Due to the many variables in each production data center, you must perform planning and testing in your own environment to validate the viability and efficacy of any solution.

NetApp HCI demonstrated impressive simplicity, performance, and QoS capabilities that outperformed a conventional HCI cluster. This enables not just greater infrastructure consolidation, but also improved application consolidation. Organizations looking for a way to provide predictable performance for business- and mission-critical databases and applications while reducing costs and complexity should take a close look at NetApp HCI.

³ Source: ESG Master Survey Results, [Converged and Hyperconverged Infrastructure Trends](#), October 2017.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.



www.esg-global.com



contact@esg-global.com



508.482.0188