

Technical Validation

# Guaranteeing Mixed-workload Performance with NetApp HCI

## Scalable, On-premises, Hyperconverged Infrastructure for Hybrid Clouds

By Tony Palmer, Senior Validation Analyst

July, 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..... 5

    Scaling Compute and Storage Resources..... 5

    ESG Testing..... 5

    Guaranteeing Performance ..... 9

    ESG Testing..... 9

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.

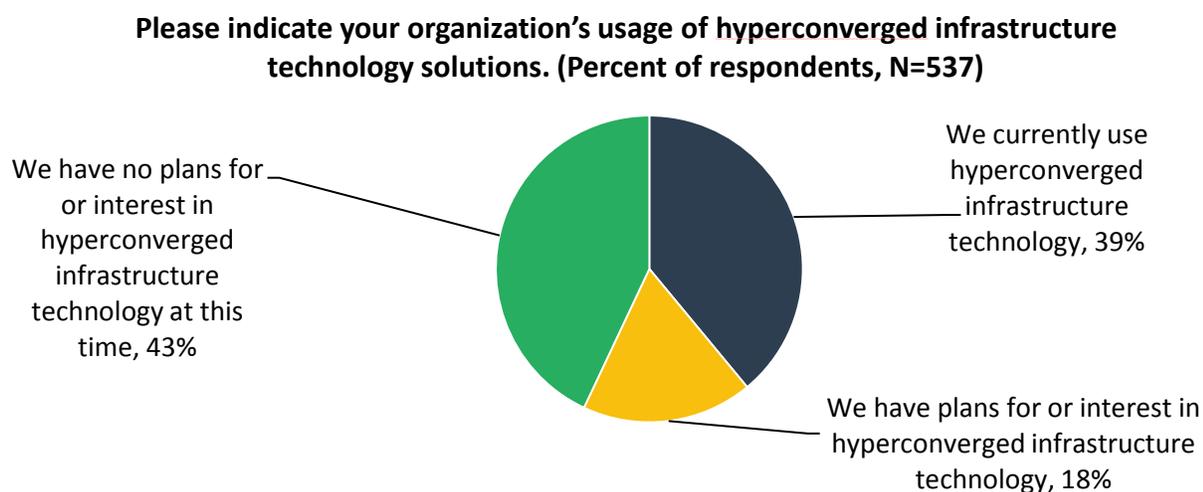
## Introduction

This ESG Technical Validation highlights the capability of NetApp HCI to deliver a guaranteed level of performance while scaling compute and storage with no effect on operations in a mixed-workload environment. Management simplicity across applications and scale is also examined.

## Background

Organizations today must be extremely agile and flexible when adding applications and virtual machines (VMs). Likewise, they must deploy business-critical production environments quickly to respond to the needs of the business. As a result, the popularity of hyperconverged infrastructure (HCI) systems has increased considerably. HCI offers a single, centrally managed solution with software-defined compute, network, and storage that is flexible, scalable, and easy to deploy. Adoption of HCI has grown significantly since coming to market, and ESG research continues to confirm the popularity of HCI; in an ESG research study, 57% of respondents reported that they use or plan to use HCI solutions. This is not surprising, given the factors driving them to consider 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%).<sup>1</sup>

**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 in order 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 over the next 12-18 months, and 93% of organizations indicated that spending on hyperconverged and converged technology would be sustained at current levels or increase in 2020.<sup>2</sup> This demonstrates that organizations view HCI as an integral part of their modernization strategy, and it is now powering mission-critical production applications.

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

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

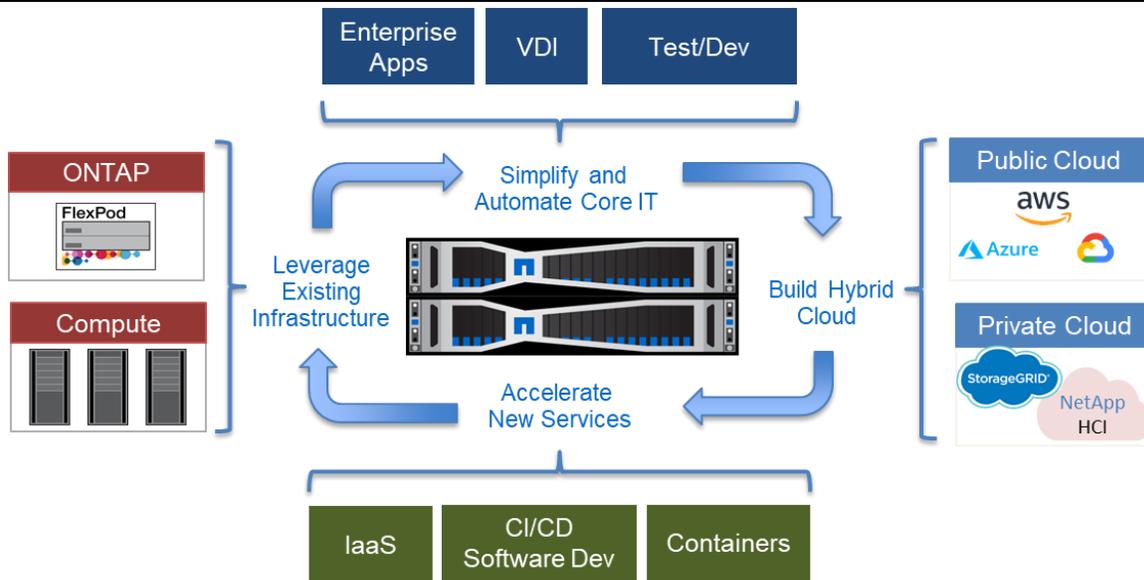
According to ESG research, adopters of HCI appreciate the simplicity of deploying and managing a tightly integrated infrastructure anchored in software-defined constructs.<sup>3</sup> Early adopters leveraged initial deployments to handle application workloads like VDI or email. Organizations looking to move mission-critical workloads traditionally reserved for three-tiered architectures or converged infrastructure (CI) solutions to HCI need to carefully consider the solution they choose. An HCI platform deployed to support tier-1 workloads must provide high IOPS and low read/write latency, *and* do so in a consistent, predictable manner. Predictable performance is critical to maximize end-user productivity across an organization. Finally, the entry-level cost is appealing to organizations that are looking to start small and grow over time.

Enter disaggregated HCI, which is designed to deliver the simplicity of first-generation HCI with more granular composability and scalability. The ability to scale compute and storage independently enables organizations to support a denser, more diverse application environment on HCI. It’s also important to note that, with this approach, organizations can reasonably expect to reduce CPU or core-based software license costs because they only have to license compute nodes, not storage nodes.

### 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

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

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 mixed-workload environment. Testing was designed to demonstrate the ability of NetApp HCI to guarantee application performance by setting minimum and maximum performance levels for each application serviced by the system. Independent scaling of storage and compute with no effect on running applications was also tested. We ran one aggressive workload and four critical workloads:

- **Aggressive workload.** This tenant generated a stressful 100% read, maximum IOPS Vdbench workload. We selected this workload because it consumed as many IOPS as were available to it and imposed maximum stress on the system.
- **Virtual desktops.** Windows 10 virtual desktops were tested with Login VSI using the knowledge worker profile, which is an intensive workload designed to stress the system smoothly, driving high CPU, RAM, and IO usage. A desktop pool of 600 knowledge workers was used for these tests.
- **SQL Server database.** Generated with HammerDB, this write-heavy workload consisted 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.
- **File services.** Run with Vdbench, this test was designed to emulate the activity of a cohort of users all accessing the same file system. A mix of reads, writes, and block sizes were used.
- **Splunk.** Splunk Enterprise was used to generate a 2TB per day ingest-rate workload that was memory- and CPU-intensive. This workload was chosen for its light storage and heavy CPU footprint to demonstrate that workloads of various characteristics can be deployed on NetApp HCI.

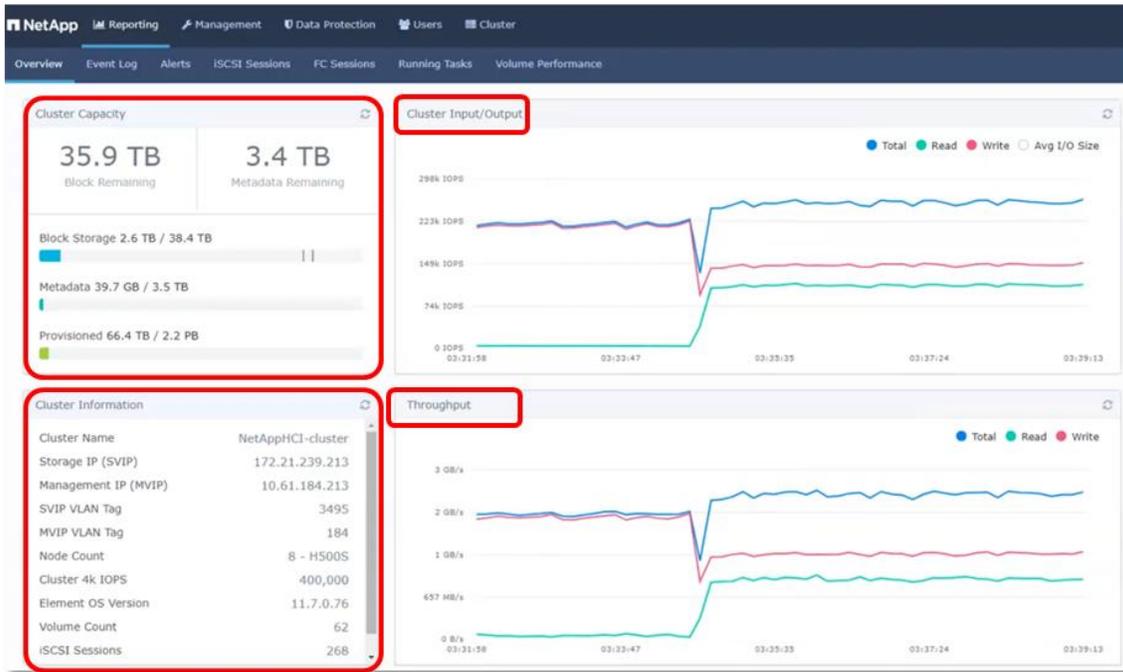
## Scaling Compute and Storage Resources

Scalability testing began with a NetApp HCI cluster with four storage nodes and two compute nodes. We ran a heavy virtual desktop workload on VMware Horizon using the Login VSI utility while we added compute and storage nodes.

## ESG Testing

First, we logged into the UI and examined the NetApp HCI Dashboard, as shown in Figure 3 and Figure 4.

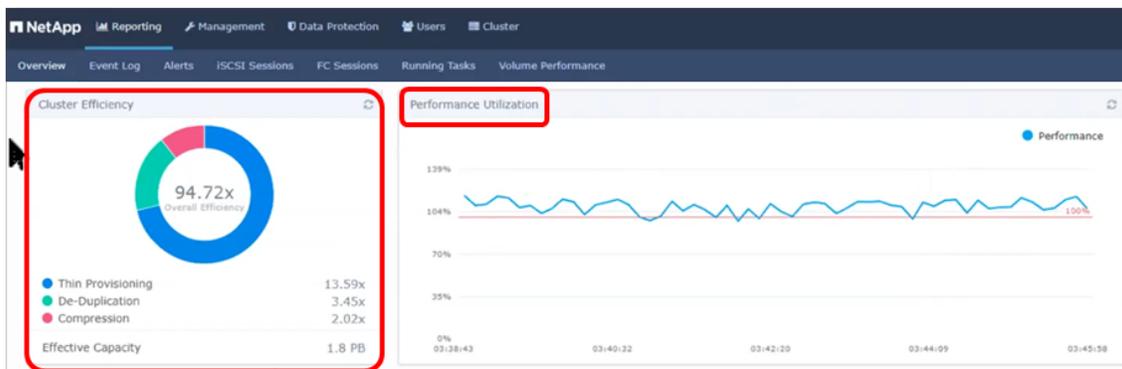
**Figure 3. The NetApp HCI Dashboard**



Source: Enterprise Strategy Group

The dashboard provides global cluster information and real-time information on storage capacity, utilization, data efficiency, and performance—both absolute IOPS/throughput and utilization based on maximum capability. Users can easily navigate to administrative tasks such as storage provisioning, data protection tasks, and cluster-related activities.

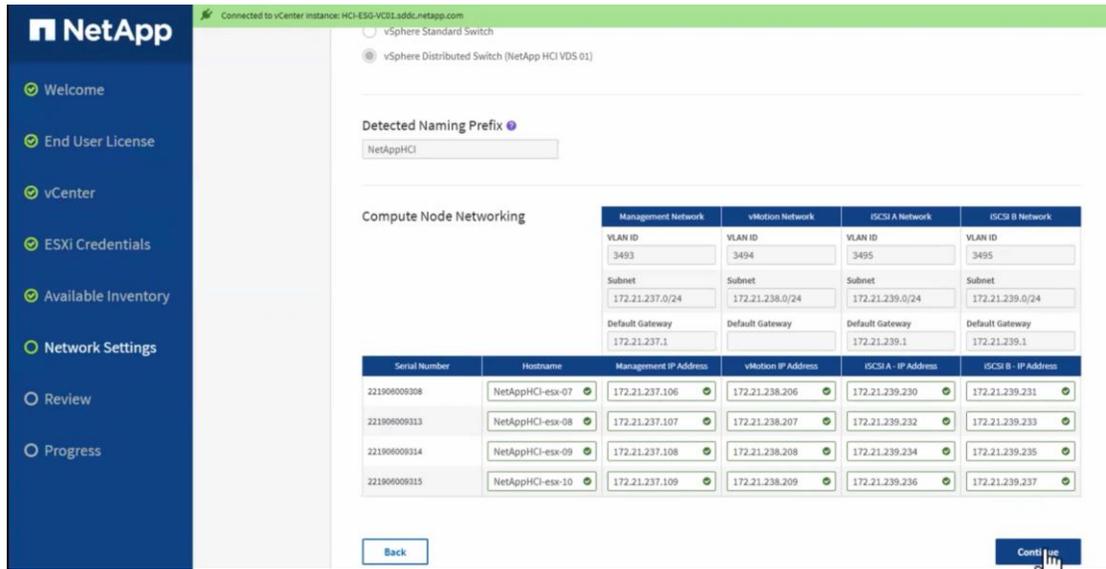
**Figure 4. Capacity Efficiency and Performance Utilization**



Source: Enterprise Strategy Group

Next, we looked at independent scalability of compute and storage resources. Our starting VDI cluster contained two compute nodes and our storage cluster contained four storage nodes. First, we added four additional compute nodes. We started a VDI workload using Login VSI configured for 150 knowledge worker desktops and ran it for the duration of the compute scale test. While the workload was running, we used the NetApp HCI UI to add four compute nodes to the cluster.

**Figure 5. Scaling Compute Resources**

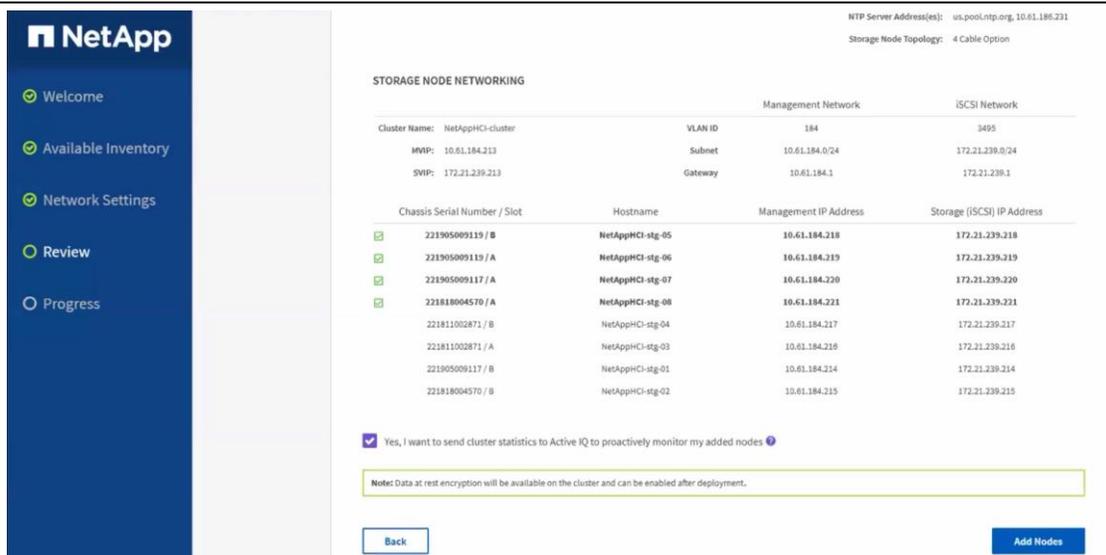


Source: Enterprise Strategy Group

The compute scaling process was simple and intuitive. We entered the ESXi credentials for our cluster, selected the already detected nodes, and configured IP network information, as seen in Figure 5. The entire process was completed in less than an hour with no effect on the running workload, and no further interaction was required. Typical administrative tasks, including but not limited to the requisite VMkernel interfaces, port groups, and attached appropriate datastores, were configured automatically. The new compute nodes were added to the cluster and the running VMs were redistributed, again completely automatically using VMware Dynamic Resource Scheduling (DRS), with no disruption to the VDI workload.

Testing the scaling of storage resources followed a similar scenario. In this test, we started the VDI workload for 200 knowledge worker desktops and ran it for the duration of the storage scale test. While the workload was running, we used the NetApp HCI UI to add four storage nodes to the cluster.

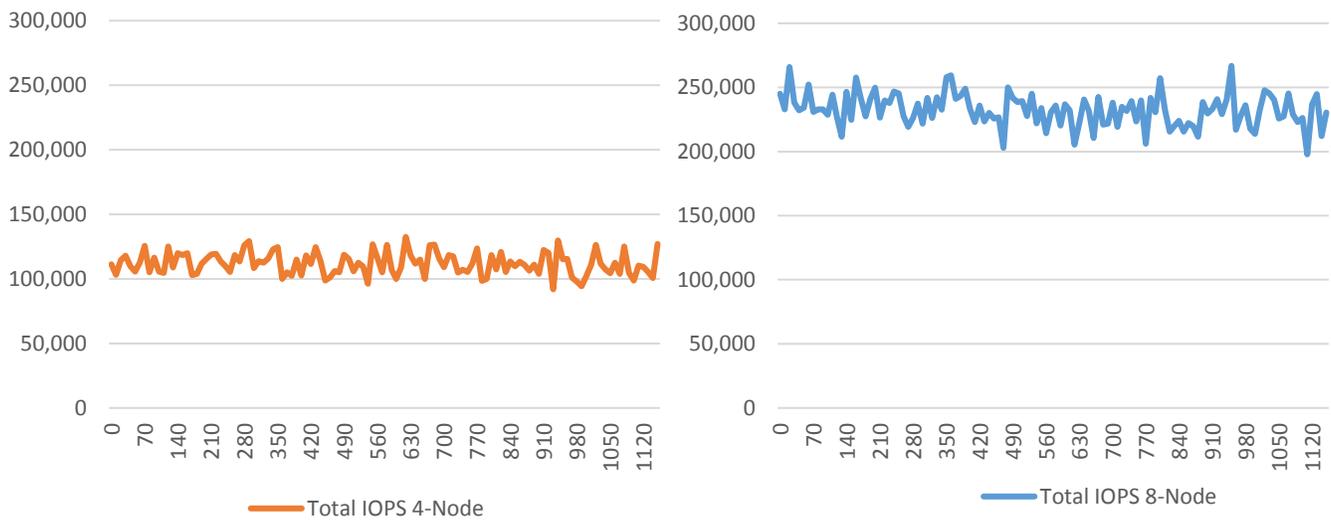
**Figure 6. Scaling Storage Resources**



Source: Enterprise Strategy Group

The storage scaling process was even simpler than adding compute nodes. We selected the four already detected new storage nodes and assigned IP network configuration, as seen in Figure 6. The entire process was completed with four clicks in 15 minutes with no effect on the running workload. No changes were needed on the compute nodes. The new storage nodes were added to the cluster, and the additional capacity was available from the single storage pool for allocation. Again, this was accomplished completely automatically and with no disruption to the VDI workload. ESG tested performance using the same configuration of servers and benchmarks before and after the storage upgrade. Before the upgrade we ran six SQL Server instances, one busy NAS server, 300 virtual desktops, and three servers running the aggressive workload. Following the scaling, we doubled all servers and benchmarks.

**Figure 7. Storage Performance—Before and After Upgrade**



Source: Enterprise Strategy Group

With four storage nodes, the system serviced an average of 111,242 total IOPS before the upgrade. After the upgrade, with eight active storage nodes, the system serviced an average of 234,442 total IOPS with our tests underway.



### Why This Matters

Organizations modernizing their data centers to support current business initiatives, including digital transformation, cloud architectures, and agile development, are challenged by overburdened personnel. According to ESG research, 34% of organizations lack IT orchestration and automation skills, 33% suffer from a lack of cloud architecture/planning skills, and 32% have a deficiency of IT architecture/planning skills.<sup>4</sup> IT organizations have discovered that delivering the simplicity, speed, accessibility, scalability, flexibility, self-service, and other benefits of HCI using traditional infrastructure and tools can be a complex and painful exercise, requiring the coordination and integration of many components.

ESG testing validated that NetApp HCI simplifies and accelerates operations and scaling of hybrid cloud infrastructures. We used automated processes from NetApp HCI to scale storage and compute in the HCI cluster independently of each other to meet business needs as they grew and changed. Scaling the cluster added performance linearly to the system with just a few clicks and minutes of actual keyboard time. There was no effect on running applications while we executed these operations, simplifying and reducing administrative workload and freeing resources for other critical IT tasks.

<sup>4</sup> Source: ESG Research Report, [2020 Technology Spending Intentions Survey](#), February 2020.

## Guaranteeing Performance

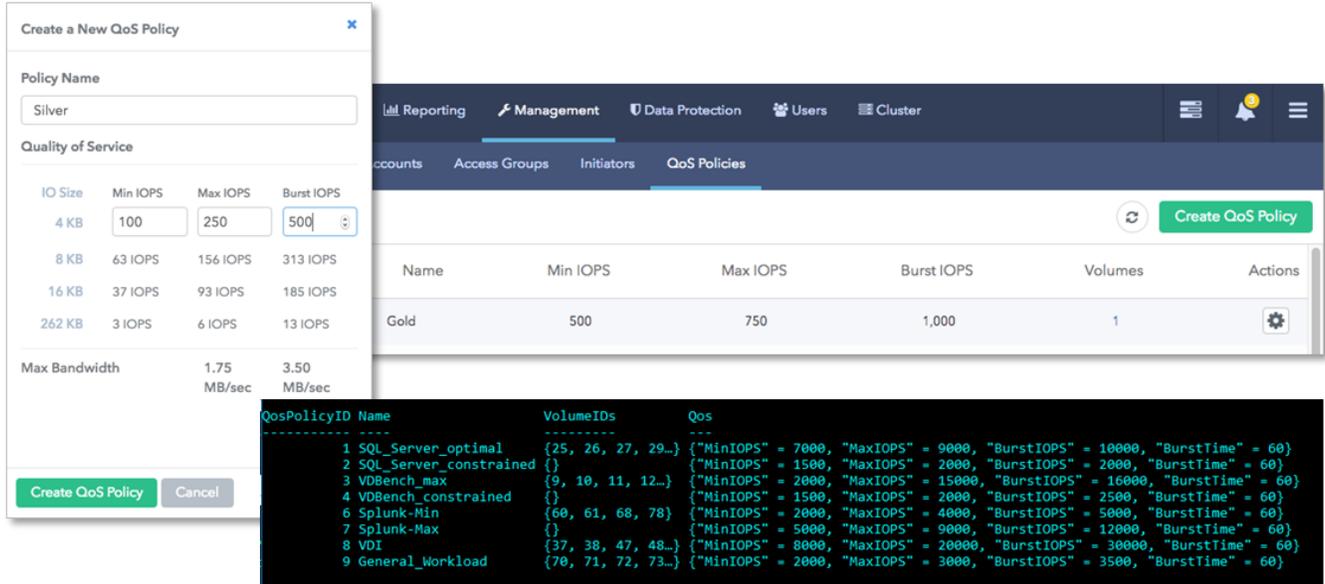
In the virtualized IT world, shared infrastructure can experience a performance imbalance from workloads sharing resources. Critical applications can be starved of storage I/O by another busy workload, resulting in unacceptable latency. NetApp HCI was designed to both limit aggressive, resource hungry 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 in order 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. This multifaceted approach protects critical app performance whether an unruly application is limited or not.

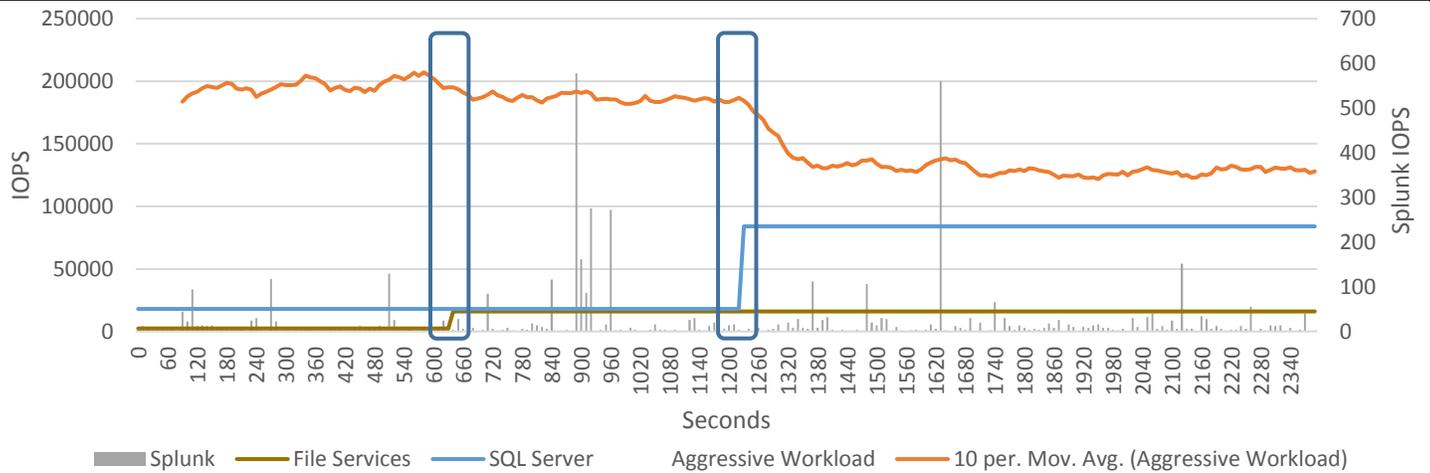
NetApp HCI QoS specifies the performance settings applied to volumes—minimum, maximum, and burst IOPS. When a volume is created with a policy, it is linked to that policy, and updates to the policy are applied to all linked volumes. We assigned all volumes associated with each workload to a different 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

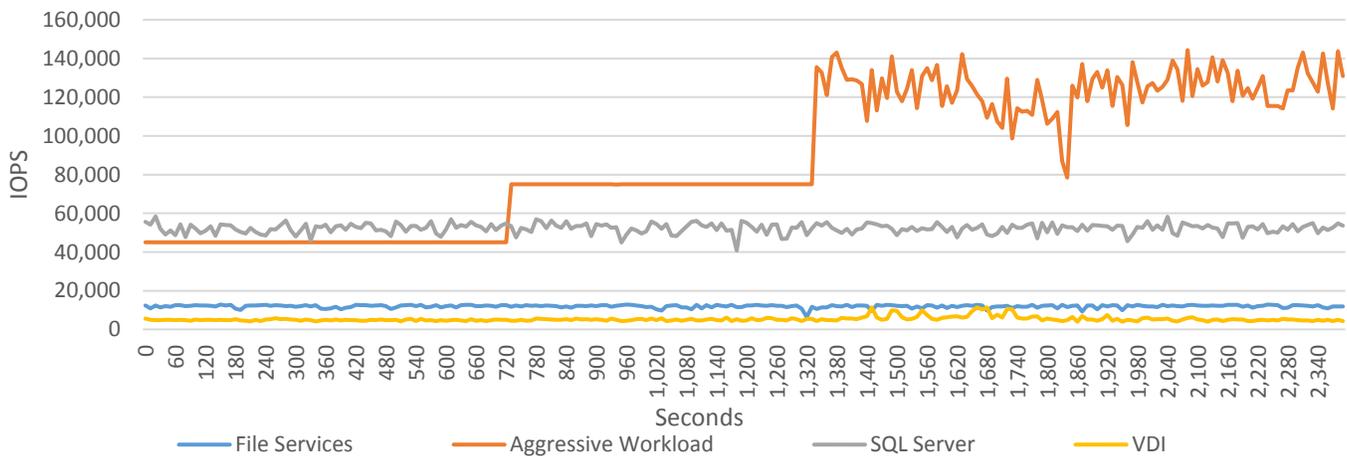
We ran all workloads simultaneously under the following conditions: Each of the critical workloads was guaranteed a specific level of performance using their linked QoS policy, and the aggressive workload was unconstrained. This phase ran for 40 minutes. Figure 9 shows the results of the first test. The chart shows the aggressive workload as a moving average of 10 data points to make it easier to follow. The Splunk workload shows actual IOPS, which were low, but Splunk Enterprise stayed responsive and processing queries throughout the test. The file services and SQL Server workloads are represented by the sum of the *MinIOPS* settings for all volumes used by those workloads.

**Figure 9. Steady State with Performance Guarantees**


Source: Enterprise Strategy Group

The file services volumes were initially set to 300 IOPS each. After 10 minutes, we increased *MinIOPS* for the file services policy to 2,000 per volume. At this point, the file services workload immediately increased to the new minimum while the unconstrained aggressive workload decreased, with the average IOPS moving from just over 200,000 to just under this value. Next, the SQL Server workload *MinIOPS* guarantee was increased from 1,500 per volume to 7,000. Again, the effect was nearly instantaneous for all 32 volumes. At this point, the aggressive workload IOPS was further reduced because the SQL Server volumes could drive more I/O.

Next, we executed another 40-minute test run. Again, we guaranteed the critical apps specific levels of performance using their linked QoS policies while we constrained the aggressive workload to 45,024 IOPS to show how precisely NetApp HCI can control performance.

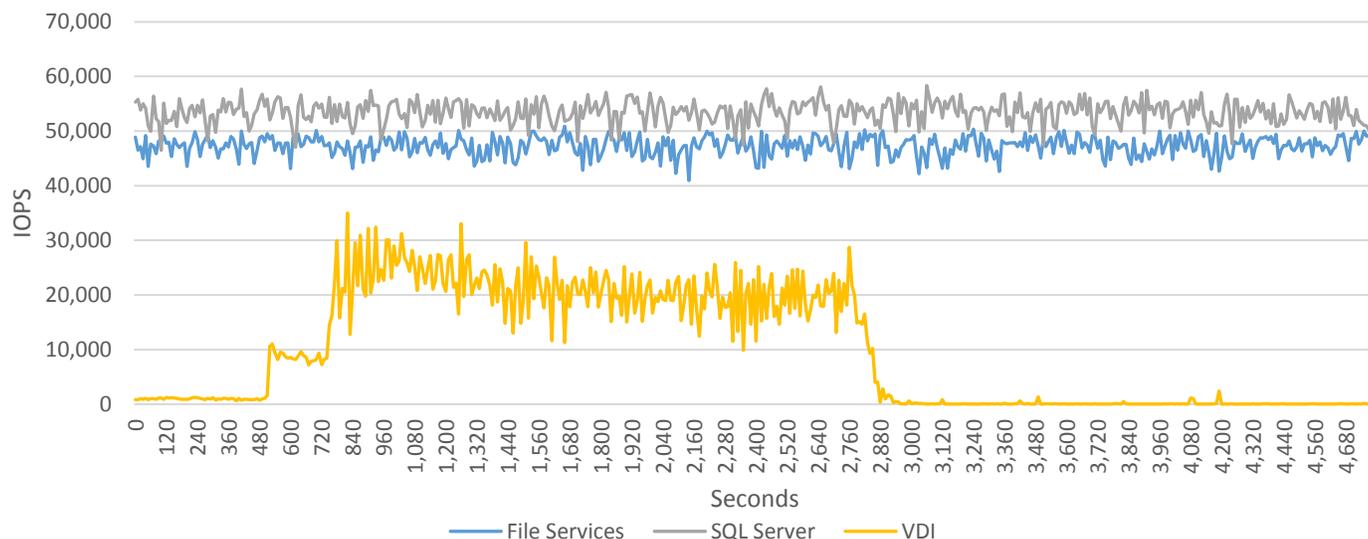
**Figure 10. Steady State with Performance Limits and Guarantees**


Source: Enterprise Strategy Group

As Figure 10 shows, for the first 12 minutes, the aggressive workload stayed at the precise level constrained by the *MaxIOPS* setting in its policy. We increased *MaxIOPS* to 75,024, and the aggressive workload immediately consumed the additional available resources, but the critical apps remained steady, unaffected by the aggressive workload. The last step of this test was to remove the *MaxIOPS* setting entirely from the aggressive workload. At this point, the unconstrained aggressive workload was consuming all the resources it could, but the critical apps were still completely unaffected thanks to their minimum performance guarantee.

Finally, a VDI boot storm of the 600 knowledge worker desktops was generated using the Login VSI tool. A boot storm occurs when many end-users boot up virtual desktops within a short timeframe, commonly overwhelming storage subsystems and degrading performance. First, we started our file services and SQL Server workloads. The file services workload was driving an average of 47,300 IOPS and the SQL Server workload was driving an average of 53,500 IOPS.

**Figure 11. VDI Boot Storm**



Source: Enterprise Strategy Group

Figure 11 shows the IOPS of each workload during the 80-minute test. IOPS for the VDI workload averaged 21,000 throughout the boot storm. It's important to note that there were no failed boots, all desktops started successfully, and the file services and SQL Server workloads were completely unaffected.

### **i** 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 applications. That leaves organizations that cannot afford performance degradation during resource contention unable to take advantage of HCI's efficiency.

ESG has validated that NetApp HCI made it simple to prioritize applications with policy-based performance. Workloads are both limited from 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 applications of varying priorities, while maintaining predictable performance. They can mix applications and types of virtual desktops, 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.

Still, 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%).<sup>5</sup>

ESG validated that NetApp HCI provides multiple benefits:

- The high performance that mission-critical, latency-sensitive applications demand, while retaining HCI's efficiency, scalability, and ease of management.
- NetApp HCI provides simple, independent resource scalability. ESG added compute and storage nodes to NetApp HCI systems independently and non-disruptively with just a few clicks. The system performed all administrative and configuration tasks automatically.
- NetApp HCI storage performance increased linearly by 108% after increasing from four storage nodes to eight.
- A VDI boot storm of 600 knowledge worker desktops completed successfully while the system was servicing more than 100,000 database and file server IOPS with no disruption to running workloads and no failed boots.
- NetApp automated, policy-based, QoS provides the right resources for every workload without constant IT intervention, managing performance based on both limits and guarantees.
- NetApp HCI enables organizations to place multiple applications—even those needing high performance—on the same HCI platform without worrying that resource contention might affect the performance of critical applications.
- NetApp Active IQ leverages AI and machine learning to provide predictive and prescriptive analytics that can help IT automate the proactive care and optimization of NetApp environments.

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.

ESG was impressed with the simplicity, performance, and QoS capabilities of NetApp HCI that enable not just infrastructure consolidation, but also application consolidation. Organizations looking for a way to reduce costs and complexity while guaranteeing service levels for critical enterprise applications would be wise to evaluate NetApp HCI.

---

<sup>5</sup> 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](http://www.esg-global.com)



[contact@esg-global.com](mailto:contact@esg-global.com)



508.482.0188