



Technical Report

## NetApp HCI QoS and Mixed Workloads

Stephen Carl, NetApp

October 2017 | TR-4632

### Abstract

This document introduces the NetApp® HCI solution to infrastructure administrators and provides important design paradigms to consider when using NetApp HCI solutions for multiple application workloads and QoS control running on the same infrastructure. The document talks about use cases that are ideally suited for HCI and discusses architecture considerations for applications running in the context of NetApp HCI. By following the guidelines in this document, you can learn how to effectively design and implement multiple application workloads on NetApp HCI.

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b> .....	<b>3</b>
1.1	Performance Guarantee .....	3
1.2	Enterprise Scale.....	3
1.3	Streamline Operations .....	3
1.4	Configuration.....	3
<b>2</b>	<b>Application Workload Consolidation</b> .....	<b>5</b>
<b>3</b>	<b>QoS Configuration</b> .....	<b>5</b>
<b>4</b>	<b>Mixed Workload Testing with QoS</b> .....	<b>8</b>
	<b>Summary</b> .....	<b>11</b>
	<b>Acknowledgements</b> .....	<b>11</b>
	<b>References</b> .....	<b>11</b>
	<b>Version History</b> .....	<b>12</b>

## LIST OF TABLES

Table 1)	NetApp HCI configuration storage nodes. ....	4
Table 2)	NetApp HCI configuration compute nodes. ....	4
Table 3)	Mongo replica set cluster VMs. ....	9
Table 4)	YCSB client server VM.....	9
Table 5)	SQL VMs.....	9
Table 6)	VDI VMs. ....	10

## LIST OF FIGURES

Figure 1)	Minimum configuration. ....	4
Figure 2)	QoS volume configuration.....	6
Figure 3)	QoS policy creation. ....	7
Figure 4)	QoS policies. ....	7
Figure 5)	VDI boot storm. ....	8
Figure 6)	Mixed workload test environment.....	9
Figure 7)	Boot storm without QoS configured for mixed workloads.....	10
Figure 8)	Boot storm after QoS is configured for mixed workloads. ....	11

# 1 Introduction

NetApp HCI is an enterprise-scale hyper converged infrastructure solution ideally suited for customers who are looking to break free from first-generation HCI limitations.

NetApp HCI customers can run multiple applications with guaranteed performance to confidently deploy resources across your entire data center. The architecture allows you to deploy your infrastructure by simplifying management and independently scale both compute and storage resources. NetApp HCI is Data Fabric ready out of the box for easy access to all your data across any public, private, or hybrid cloud. By moving to NetApp HCI, IT organizations can transform their data center, driving operational efficiencies and reducing costs.

Data Fabric is a software-defined approach from NetApp for data management that enables businesses to connect disparate data management and storage resources. NetApp HCI can streamline data management between on-premises and cloud storage for enhanced data portability, visibility, and protection.

## 1.1 Performance Guarantee

A common challenge for a data center is delivering predictable performance, complicated even more by running multiple applications sharing the same infrastructure. An application interfering with other applications creates performance degradation, causing IT administrators to spend valuable time troubleshooting the environment. Mainstream applications, such as virtual desktop infrastructure (VDI) and database applications, have unique I/O patterns that can, during normal operations, affect one another's performance when deployed in a shared environment. NetApp's HCI quality of service (QoS) feature allows fine-grained control of performance for every application, eliminating noisy neighbors, meeting unique performance needs, providing higher utilization of infrastructure, and satisfying performance SLAs. The storage architecture which is part of the NetApp HCI solution eliminates performance variance in the context of data locality as the data is distributed across all the nodes in the HCI cluster.

## 1.2 Enterprise Scale

Unlike previous generations of HCI which have fixed resource ratios, NetApp HCI scales compute and storage resources independently. Independent scaling avoids costly and inefficient overprovisioning and simplifies capacity and performance planning. Running on innovative SolidFire® technology and delivered on a NetApp designed architecture, NetApp HCI is an enterprise-scale hyper converged infrastructure solution. NetApp HCI comes in 2RU 4-node building blocks (chassis) in mix-and-match small, medium, and large storage and compute configurations and can be scaled to allow you to rapidly meet changing business needs and scale on your terms.


## 1.3 Streamline Operations

A common goal of IT organizations is to automate all routine tasks, eliminating the risk of user errors associated with manual operations, allowing valuable resources to be focused on higher value priorities that drive business efficiencies. The NetApp Deployment Engine (NDE) streamlines day 0 installation from hours to minutes, while simple centralized management through the vCenter plug-in gives you full control of managing your entire infrastructure through an intuitive user interface. A robust suite of APIs enables additional seamless integration into higher-level management, orchestration, backup, and disaster recovery tools.

## 1.4 Configuration

NetApp HCI is available with configuration options of small, medium, or large for both compute and storage. The nodes are similar to a small blade that sits inside a chassis. A minimum starting configuration must have four storage nodes and two compute nodes.

Figure 1) Minimum configuration.



From the configuration information in Table 1, each storage node can deploy from 5TB to 44TB of effective capacity. From a compute node perspective, shown in Table 2, 16 to 36 CPU cores and 384GB to 768GB of RAM are available.

Table 1) NetApp HCI configuration storage nodes.

	Small	Medium	Large
RU	1RU, half-width	1RU, half-width	1RU, half-width
CPU	E5-2620 v4; 8C at 2.1GHz	E5-2650 v4; 12C at 2.2GHz	E5-2695 v4; 18C at 2.1GHz
Boot device	1 x 240GB MLC	1 x 240GB MLC	1 x 240GB MLC
Base networking	(2) 25/10GbE SFP28/SFP+ (2) 1GbE RJ45	(2) 25/10GbE SFP28/SFP+ (2) 1GbE RJ45	(2) 25/10GbE SFP28/SFP+ (2) 1GbE RJ45
SSD	(6) 480GB	(6) 960GB	(6) 1.9TB
Effective block capacity	5.5TB–11TB	11TB–22TB	22TB–44TB

Table 2) NetApp HCI configuration compute nodes.

	Small	Medium	Large
RU	1RU, half-width	1RU, half-width	1RU, half-width
Cores for VMs	16	24	36
CPU	E5-2620 v4; 8C at 2.1GHz	E5-2650 v4; 12C 2.2GHz	E5-2695 v4; 18C at 2.1GHz
Memory	384GB	512GB	768GB
Boot device	(1) 240GB MLC	(1) 240GB MLC	(1) 240GB MLC
Base networking	(4) 25/10GbE SFP28/SFP+ (2) 1GbE RJ45	(4) 25/10GbE SFP28/SFP+ (2) 1GbE RJ45	(4) 25/10GbE SFP28/SFP+ (2) 1GbE RJ45

As an example, for a minimum size starting solution, a configuration with two small compute and four small storage nodes would have 32 cores with 768GB of memory and between 22TB to 44TB of effective capacity. As requirements change, you can add more compute or storage nodes of any size to the chassis independently of each other. This flexibility of adding only compute or only storage nodes enables unique scalability options for building an efficient and agile cloud in your data center for a variety of use cases.

## 2 Application Workload Consolidation

Consolidate all of your workloads, avoiding the need for you to scale in ways that strand resources and throttle the performance required by next generation data center applications. NetApp HCI provides the solution for IT predictability challenges with unique QoS limits that allow the granular control of every application, eliminating noisy neighbors, meeting unique performance needs, and satisfying all performance SLAs.

The per-volume QoS control of HCI storage nodes helps individual virtual machines (VMs) and applications meet I/O throughput requirements without being affected by other applications running in parallel. With QoS and data reduction efficiencies, higher application density with the shared storage infrastructure can be achieved.

VMware administrators have full control of each storage volume on which the application data resides and can perform all maintenance operations, including setting the QoS for each data volume copy through the vCenter SolidFire plug-in. Administrators can also use NetApp HCI REST APIs to achieve full automation and make storage management simple and easy.

### Granular Control

The ability to dynamically allocate, manage, and guarantee performance independently of capacity is unique to the NetApp HCI product. This is achieved through the ability to define minimum, maximum, and burst levels for each and every application or volume with the QoS feature.

Segmenting the performance of individual workloads eradicates the concept of the noisy neighbor (where one workload steps on top of another, causing a performance issue that reverberates across the infrastructure).

The combined effect of tying performance settings to individual VMs, as well as applying automation constructs around storage-based policy management, is one of the most powerful steps toward an automated, policy-driven infrastructure. The guaranteed performance offered by NetApp HCI lets you consolidate mixed workloads, deliver predictable performance, and provide granular control at the VM level.

## 3 QoS Configuration

Configuring hosts for QoS offers granular control for different applications and workloads that can compete for resources. The uniqueness of NetApp HCI QoS lies in its ability to provision both performance and capacity resources independently; guarantee both per workload; and allow admins to know, monitor, trend, manage, and adjust those allocations in real time.

Configuring a QoS policy for a volume is done during volume creation or can be changed later. The NetApp SolidFire management UI allows the user to define minimum, maximum, and burst levels for IOPS application volumes.

QoS settings:

- **Minimum IOPS.** The minimum number of IOPS guaranteed for the volume.
- **Maximum IOPS.** The maximum number of IOPS allowed for the volume.

- **Burst IOPS.** The maximum number of IOPS allowed over a short period of time for the volume.

Figure 2 shows how to configure storage volume QoS to support a multitenant environment for multiple virtual machines and applications running separate workloads on a NetApp HCI infrastructure.

To configure or edit QoS on a volume, follow these steps:

1. Log in to the NetApp SolidFire management UI from the vCenter UI.
2. Click Volume → Edit Volume on the far right.

Figure 2) QoS volume configuration.

**MongDAT3 - Edit Volume**

Volume name: MongDAT3  
 Volume ID: 12  
 IQN: iqn.2010-01.com.solidfire:w4fm.mongdat3.12  
 Account: NetApp-HCI  
 Size: 1000.00 GB  
 512e: Yes

---

**Volume Details**

Volume size \* 1000 GB  
 Access \* Read/Write  
 Account NetApp-HCI

---

**Quality of Service**

I/O Size	Min IOPS	Max IOPS	Burst IOPS
4 KB	15000	100000	100000
8 KB	9375 IOPS	62500 IOPS	62500 IOPS
16 KB	5556 IOPS	37037 IOPS	37037 IOPS
256 KB	385 IOPS	2564 IOPS	2564 IOPS

Max Bandwidth 699.05 MB/sec 699.05 MB/sec

Reset Default QoS

OK Cancel

QoS configuration is also configurable by creating a QoS policy in the SolidFire UI. Figure 3) QoS policy creation.

shows how a QoS policy is created.

3. Log in to the NetApp SolidFire UI.
4. Click the Management tab → QoS Policies.
5. Click → Create QoS Policy.

Figure 3) QoS policy creation.

### Create a New QoS Policy ✕

---

**Policy Name**

**Quality of Service**

IO Size	Min IOPS	Max IOPS	Burst IOPS
4 KB	10000	35000	40000 ▾
8 KB	6250 IOPS	21875 IOPS	25000 IOPS
16 KB	3704 IOPS	12963 IOPS	14815 IOPS
262 KB	256 IOPS	897 IOPS	1026 IOPS
Max Bandwidth		244.67 MB/sec	279.62 MB/sec

Figure 4) QoS policies.

shows a list of QoS policies used in this document's testing.

Figure 4) QoS policies.

### Quality of Service

---

Policy

Select a QoS Policy

- mongo-production**  
 Min: 10000, Max: 35000, Burst: 40000
- mongo-QoS  
 Min: 10000, Max: 25000, Burst: 26000
- sql-QoS  
 Min: 15000, Max: 30000, Burst: 32000
- VDI  
 Min: 5000, Max: 5000, Burst: 5000


When the QoS minimum, maximum, and burst IOPS are configured, volume protection of application performance is enforced during conditions when workloads can require a high spike of needed resources. For example, a VDI boot storm of a significant number of end users booting up within a very narrow time

frame can drown out other workloads. This behavior has necessitated dedicating resources to VDI workloads. Figure 5) VDI boot storm.

Figure shows a typical system performance spike when a VDI boot storm occurs.

Figure 5) VDI boot storm.

Figure 5) VDI boot storm.



Configuring volumes with QoS policy granular control can prevent overallocation of resources to any specific type of workload.


## 4 Mixed Workload Testing with QoS

The testing done for this document used an environment of several VMs for VDI hosts, MongoDB clusters, and SQL databases. Each database ran a common workload, while the VDI boot storm was simulated by powering on several datastores containing several VDI hosts. This simulation is similar to several end users powering up systems in the same time frame.

Figure 6) Mixed workload test environment.

is the logical representation of the test environment of different applications configured in the NetApp HCI cluster.

Figure 6) Mixed workload test environment.



The following tables show the VMs configured for the application workloads.

Table 3) Mongo replica set cluster VMs.

MongoDB Cluster	Qty	MongoDB Version	Type	Cores/vCPUs	RAM	OS
Cluster 1	3	3.4	VMware	8	32GB	Red Hat Linux
Cluster 2	3	3.4	VMware	8	32GB	Red Hat Linux

Table 4) YCSB client server VM.

YCSB Client	Qty	YCSB version	Type	Cores/vCPUs	RAM	OS
YCSB	1	0.12.0	VMware	8	32GB	Red Hat Linux

Table 5) SQL VMs.

MongoDB Cluster	Qty	Type	Cores/vCPUs	RAM	OS
Hosts	5	VMware	8	32GB	Red Hat Linux


Table 6) VDI VMs.

VM Type	Qty	Type	vCPUs	RAM	OS
Virtual desktops	400	VMware linked clones	2	4GB	Windows 10

The volume view of this VDI boot storm scenario without QoS configured in Figure 7) Boot storm without QoS configured for mixed workloads.

shows the performance and resources impact, including other application workloads.


Figure 7) Boot storm without QoS configured for mixed workloads.



After configuration with the NetApp HCI QoS controls to fine tune IOPS limits of minimum (guaranteed), maximum, and burst settings, the same VDI boot storm is tested again, as shown in Figure 8) Boot storm after QoS is configured for mixed workloads.

. The high spike in VDI consumed resources for performance is prevented, and the application performances of other application workloads are protected and retain their QoS-enabled performance. The noisy neighbor concept on affecting mixed application workloads performance is eliminated by QoS.

Figure 8) Boot storm after QoS is configured for mixed workloads.



## Summary

Data centers are moving away from dedicated platforms and are trying to avoid overprovisioning to increase efficiency and reduce cost. NetApp HCI provides a unique solution with QoS limits, allowing the granular control of every application, eliminating noisy neighbors, and satisfying all performance SLAs. The underlying storage nodes of NetApp HCI use all-flash media coupled with capacity thin provisioning and in-line data reduction features. These features yield significant efficiency and agility when deploying applications and help the business in consolidating workloads with confidence. This benefits system planners and administrators when deploying and maintaining database applications such as MongoDB. For additional information, see [NetApp SolidFire](#).

## Acknowledgements

The authors would like to thank Chris Gebhardt, Chris Merz, Bobby Oommen, and Aaron Patten for their contributions to this technical report.

## References

NetApp HCI Datasheet  
<http://www.netapp.com/us/media/ds-3881.pdf>

## Version History

Version	Date	Document Version History
Version 1.0	October 2017	Initial document creation

## Copyright Information

Copyright © 2017 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.