



Technical Report

Performance characterization of NetApp Cloud Volumes ONTAP for Google Cloud

Danny Tzidony and Michele Pardini, NetApp Hybrid Cloud Services Group
January 2023 | TR-4816

Abstract

This technical report examines the performance of NetApp® Cloud Volumes ONTAP® for Google Cloud for application workloads. NetApp partners, customers, and employees should use the presented information to make informed decisions about which workloads are appropriate for Cloud Volumes ONTAP.

TABLE OF CONTENTS

Introduction	4
Executive summary	4
Measurement of storage performance	4
Test methodology	4
Improvements in ONTAP 9.11.1 storage efficiencies	5
Cloud Volumes ONTAP configuration	6
Workloads.....	7
Test environment diagram.....	7
Performance assessment results	8
Transactional/OLTP.....	8
Streaming reads	10
Streaming writes.....	12
Mixed reads/writes.....	13
Random reads	16
Random writes.....	18
Acknowledgements	20
Where to find additional information	20
Version history	21

LIST OF TABLES

Table 1) Google Compute Engine machine types, performance limits, and selected storage configuration.	6
Table 2) Transactional/OLTP workload results for single-node configuration.....	9
Table 3) Transactional/OLTP workload results for a HA configuration.	10
Table 4) Streaming reads workload results for single-node configuration.	11
Table 5) Streaming reads workload results for HA configuration.....	12
Table 6) Streaming writes workload results for single-node configuration.....	12
Table 7) Streaming writes workload results details for HA configuration.	13
Table 8) Mixed reads/writes 16KB result details for single-node configuration.....	14
Table 9) Mixed reads/writes 16KB result details for HA configuration.	15
Table 10) Mixed reads/writes 32KB result details for single-node configuration.....	16
Table 11) Mixed reads/writes 32KB result details for HA configuration.	16
Table 12) Random reads 8KB result details for single-node configuration.....	17
Table 13) Random reads 8KB result details for HA configuration.	18
Table 14) Random writes 8KB result details for single-node configuration.....	19

Table 15) Random writes 8KB result details for HA configuration	20
---	----

LIST OF FIGURES

Figure 1) Test environment diagram.	8
Figure 2) Transactional/OLTP workload IOPS for single-node configuration.	9
Figure 3) Transactional/OLTP workload IOPS for a high-availability (HA) configuration.	10
Figure 4) Streaming reads throughput for single-node configuration.....	11
Figure 5) Streaming reads throughput for HA configuration.	11
Figure 6) Streaming writes throughput for single-node configuration.	12
Figure 7) Streaming writes throughput for HA configuration.....	13
Figure 8) Mixed reads/writes 16KB throughput for single-node configuration.	14
Figure 9) Mixed reads/writes 16KB throughput for HA configuration.....	15
Figure 10) Mixed reads/writes 32KB throughput for single-node configuration.	15
Figure 11) Mixed reads/writes 32KB throughput for HA configuration.....	16
Figure 12) Random reads 8KB IOPS for single-node configuration.	17
Figure 13) Random reads 8KB IOPS for HA configuration.....	18
Figure 14) Random writes 8KB IOPS for single-node configuration.....	19
Figure 15) Random writes 8KB IOPS for HA configuration.	20

Introduction

NetApp has a long history of providing leading-edge, high-performance, feature-rich data management technologies and solutions that provide the agility and mobility that organizations need in the digital transformation and cloud era. The NetApp Cloud Volumes ONTAP software-only storage solution is part of the ONTAP family of products that delivers an organization's data anywhere that it is needed across the hybrid multicloud.

NetApp Cloud Volumes ONTAP is an enterprise-grade data management solution that is deployed on a public cloud infrastructure. It consumes native cloud compute services and manages native cloud block and object storage. Today, Cloud Volumes ONTAP can be deployed in any of the "big three" cloud providers: Google Cloud, Amazon Web Services (AWS), and Azure. Through its unified storage access capabilities, storage efficiency features, and enhanced data protection and security services, Cloud Volumes ONTAP delivers a powerful solution for any enterprise workload. It is optimal for workloads such as databases, DevOps, virtual desktop infrastructure (VDI), Kubernetes, file services, disaster recovery, backup and archiving, and more.

This technical report describes the results of performance assessments that were conducted with Cloud Volumes ONTAP for Google Cloud. When deployed in Google Cloud, Cloud Volumes ONTAP uses Google Compute Engine, a scalable high-performance virtual machine (VM), and Google Persistent disk, which is durable and high-performance block storage. It also uses Google Cloud Storage, which provides unified, price/performance-optimized object storage.

Cloud Volumes ONTAP is available in two pricing models: pay as you go (PayGo) and bring your own license (BYOL). For the PayGo model, four types of license packages are available: Freemium, Essential, Professional and Optimized. Each license package provides different compute and storage capacity options.

Executive summary

NetApp performed this study to showcase the storage performance and benefits of NetApp Cloud Volumes ONTAP. Understanding the performance characteristics of Cloud Volumes ONTAP is critical to set expectations and to help stakeholders make educated decisions about which of their workloads Cloud Volumes ONTAP can handle and what configuration is required.

The results are presented in section 0, Performance assessment results. The results provide some basic understanding of the Cloud Volumes ONTAP performance capabilities for Google Cloud and its dependency on the underlying infrastructure and the services that are used.

The chosen workloads for this study were transactional/OLTP, streaming reads, streaming writes, mixed reads/writes and random reads/writes. Each of the selected workloads is characterized by different mixture and access patterns of I/O operations.

Measurement of storage performance

NetApp performed these assessments to measure the performance of single-node and high-availability (HA) NetApp Cloud Volumes ONTAP 9.11.1 in different application workloads. This section describes the methodology that NetApp used to assess Cloud Volumes ONTAP performance and the configuration that was used, as well as the workloads that were tested.

Test methodology

For this study, NetApp used Flexible I/O Tester (Fio version fio-3.30-63-g66087), a tool that can simulate any given workload. Fio generates a number of threads or processes, carrying out the type of I/O

combination and patterns that the user specifies to generate the desired workloads. For more information about Fio, see the [Fio page on GitHub](#).

To get the most out of Cloud Volumes ONTAP, NetApp used multiple Fio clients in the various assessments. The I/O workloads were generated by multiple n2-standard-8 compute engines (each with eight virtual CPUs [vCPUs] and 30GB of memory) running CentOS Linux 8 (for a total of six).

All tests were performed with NFSv4, which is commonly used in a wide variety of workloads. NFS is a distributed file system, based on a client-server architecture, that enables users to access files and directories on a remote file system, much like the process to access local storage. The NFS protocol provides file-level access to logical storage resources, such as NetApp FlexVol® volumes, on Cloud Volumes ONTAP. A FlexVol volume is an allocation of storage space that serves as a data container. A FlexVol volume is the basic storage unit that can be mounted by one or more NFS clients, providing access to users within the limits of the granted permissions.

In this study, each client mounted a separate FlexVol volume. The goal was to achieve the best performance metrics possible for IOPS, throughput (the amount of data payload), and latency (the response time). All the results that are presented later in this report are the aggregated sum from the clients' side.

All the Compute Engine instances that NetApp used in the study, including the instances that Cloud Volumes ONTAP used, were in the same availability zone, Virtual Private Cloud, and subnet.

Improvements in ONTAP 9.11.1 storage efficiencies

ONTAP 9.11.1 introduces a new storage efficiency capability, also known as temperature sensitive storage efficiency, enabling Cloud Volumes ONTAP to efficiently perform inline storage efficiencies such as compression and compaction. Temperature sensitive storage efficiency can lead to significant savings in capacity and performance enhancements, which depend on the compression rate of the data.

- **Inline compression.** By default, ONTAP 9.11.1 uses a compression group with 8K block size to determine how to compress new data. The storage efficiency algorithm determines whether to apply compression or not, if at least 50% of the data can be saved. Meaning, for 8K, at least 4K must be saved, so the compress rate is lower or equal to 50%. If for example, a workload has a compress rate of less than 50%, there is little to no saving.
- **The importance of inline storage efficiencies in Cloud Volumes ONTAP.** Cloud Volumes ONTAP has limitations deriving from the Google Cloud infrastructure services it utilizes. Two notable limitations stem from Compute Engine limitations and persistent disks limitations. These limitations influence the handling capability of input/output operations per second (IOPS) and the amount of throughput (MB/s).
- **The effect of these limitations can be mitigated by inline storage efficiencies if applied before the data is written to disk.** This means that Cloud Volumes ONTAP can accept and handle more IOPS and throughput from the same clients. It is important to note that the limits themselves cannot be exceeded, but with inline storage efficiencies more clients' workload can be handled by Cloud Volumes ONTAP.

In addition to inline storage efficiencies, several performance enhancements were introduced in recent Cloud Volumes ONTAP releases (starting 9.9) that minimize the number of operations to enhance the write performance.

Therefore, by using ONTAP 9.11.1 inline storage efficiencies allow users to:

- Reduce TCO by running the same workload on smaller instances or add more workload to an existing solution.
- Get better write response times, up to the optimal point.

Note: The above assumes that most, if not 100% of the workload can be 50%+ reducible. If only a fraction of the workload can be reducible then improvements would be negligible, resembling

the performance of previous Cloud Volumes ONTAP releases. Additionally, workloads dominated by reads, or 100% reads, might not get performance benefit from inline storage efficiencies.

Cloud Volumes ONTAP configuration

For this study, Cloud Volumes ONTAP working environments were configured as a single-node or a high-availability (HA) system. The performance of Cloud Volumes ONTAP depends on the underlying cloud services that are used (the Compute Engine machine type and the type and size of the Persistent Disk). Therefore, NetApp tested Cloud Volumes ONTAP by using different Compute Engine machine types and Persistent disk configurations. The Compute Engine machine types not only differ in their specifications, but they also incur different costs by Google Cloud.

NetApp used the following Compute Engine machine types:

- n2-standard-8
- n2-standard-16
- n2-standard-32

All the VM instances were configured with SSD persistent disks. SSD persistent disks are backed by SSDs and provide predictable performance, having been designed for single-digit, millisecond (ms) latencies. Because they balance price/performance for a wide variety of transactional workloads, the use of SSD persistent disks is preferable. Cloud Volumes ONTAP supports the use of Zonal Standard persistent disks as well, which are backed by HDDs and provide an efficient and economical way to handle sequential read/write operations. For more information about Google Cloud persistent disks and limits, visit the [Google Cloud storage options page](#).

Table 1 displays the Compute Engine machine types, performance limits, and storage configuration that NetApp used in this study.

Table 1) Google Compute Engine machine types, performance limits, and selected storage configuration.

Machine type	vCPU	Memory (GiB)	Sustained random IOPS		Sustained throughput (MBps)		SSD persistent disk
			Read	Write	Read	Write	
n2-standard-8	4	16	15,000	15,000	240	240	4 PD-SSD x 2TB per aggr*
n2-standard-16	8	32	15,000	15,000	800	800	4 PD-SSD x 2TB per aggr*
n2-standard-32	32	128	50,000	50,000	1,200	1,200	4 PD-SSD x 2TB per aggr*

Note: For more information about the Google Cloud performance limitations, visit the Google Cloud [Optimizing persistent disk performance page](#).

*= HA tested in A/A mode, so 2 aggregates with 4 PD-SSD * 2TB each

The Cloud Volumes ONTAP write speed, which affects write performance, was set to `Normal`. When the write speed is set to `Normal`, data is first written to persistent disk storage and then an acknowledgment is sent back to the client. This process eliminates the chance of data loss if an unplanned outage occurs.

When the write speed is set to `High`, data is buffered in memory and an acknowledgment is sent back to the client before the data is committed to persistent disk storage. A high write speed enhances write performance but also increases the potential for data loss if an unplanned outage occurs. Users should enable a high write speed only if enhanced write performance is required and if data loss can either be

tolerated or be handled by the application. For example, data loss from an unplanned outage can be tolerated, or the application (such as Microsoft SQL Server for Always On availability groups or MySQL replication) can handle data loss.

As of June 2022, high write speed is not supported on a HA configuration.

Workloads

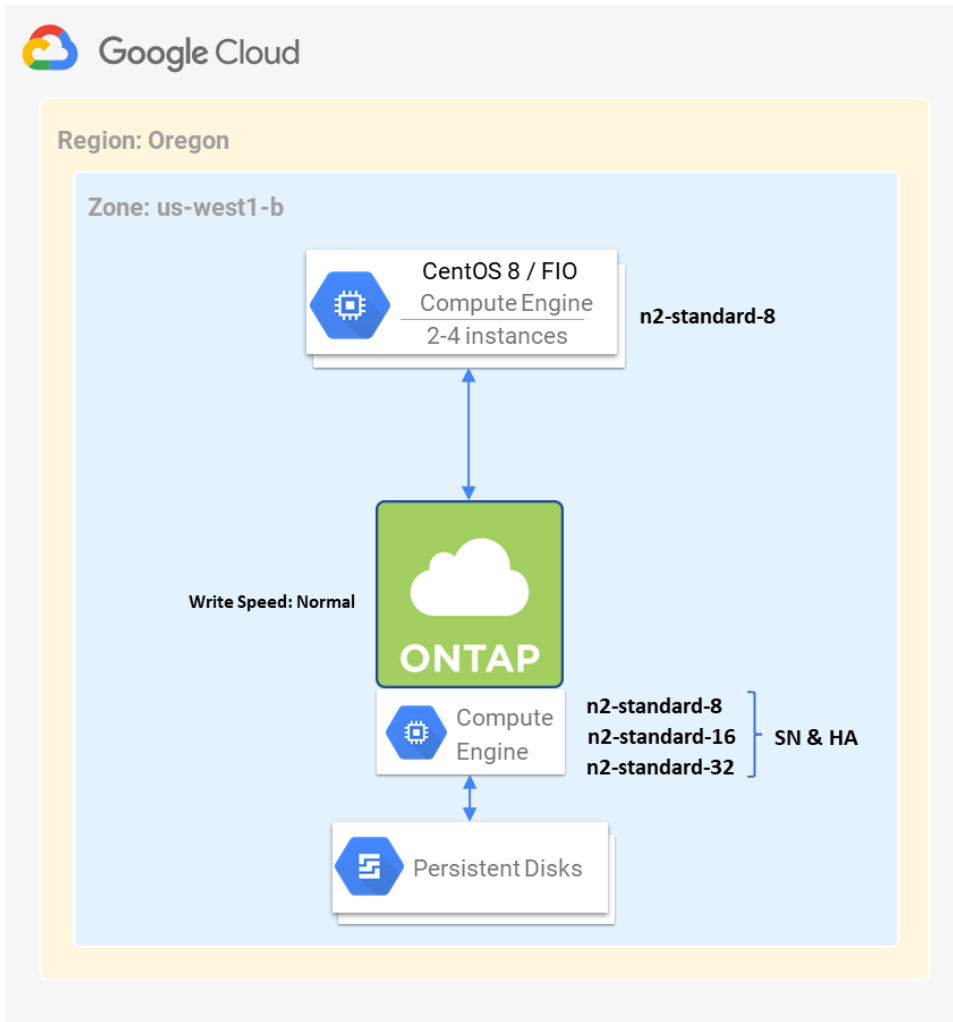
In this study, four different workloads were configured and tested. Each tested workload consisted of a unique I/O mixture and access pattern that simulated workloads that are used in widely deployed applications:

- **Transactional/OLTP** (8KB block size, 80% reads, and 100% random access I/O). A transactional workload is typically identified by a database that receives numerous requests for data and multiple changes to this data from several users over time. These modifications are called transactions. The transactions are small and primarily random in nature with high concurrency. This type of workload is generated by transactional database applications such as SAP, Oracle, and SQL Server and MySQL databases.
- **Streaming/sequential reads** (64KB block size, 100% reads, and 100% sequential access I/O). In a streaming reads workload, only read operations are performed (and they are concurrent, large, contiguous requests). This workload is typically generated by applications such as media servers (for example, video on demand) and virtual tape libraries.
- **Streaming/sequential writes** (64KB block size, 100% writes, and 100% sequential access I/O). In a streaming writes workload, only write operations are performed (and they are concurrent, large, contiguous requests). This workload is typically generated by applications such as media capture, virtual tape libraries, medical imaging, backup and archiving, and video surveillance.
- **Mixed reads/writes** (32 KB and 16KB block size, 50% reads, and 100% random access I/O). Many applications generate a mix of read/write operations that differ in the reads versus writes percentage, block size, and random versus sequential percentage and that typically require high throughput and low latency. A mix of 50% reads and 50% writes is a common starting point for measuring. That kind of workload can be found, for example, in a VDI.
- **Random reads** (8KB block size, 100% reads, and 100% random access I/O). One of the most required workloads for sizing purposes. In addition, this workload pattern is seen in databases sporadically.
- **Random writes** (8KB block size, 100% reads, and 100% random access I/O). One of the most required workloads for sizing purposes. In addition, this workload pattern is seen in databases sporadically.

Test environment diagram

Figure 1 displays the environment and Google Cloud resources that NetApp used for the performance assessments.

Figure 1) Test environment diagram.



Performance assessment results

This section describes the values of the performance metrics that were obtained in the process of assessing NetApp Cloud Volumes ONTAP.

Transactional/OLTP

The significant metrics for this type of workload are IOPS, as shown in Figure 2 and as described in Table 2

Figure 2) Transactional/OLTP workload IOPS for single-node configuration.

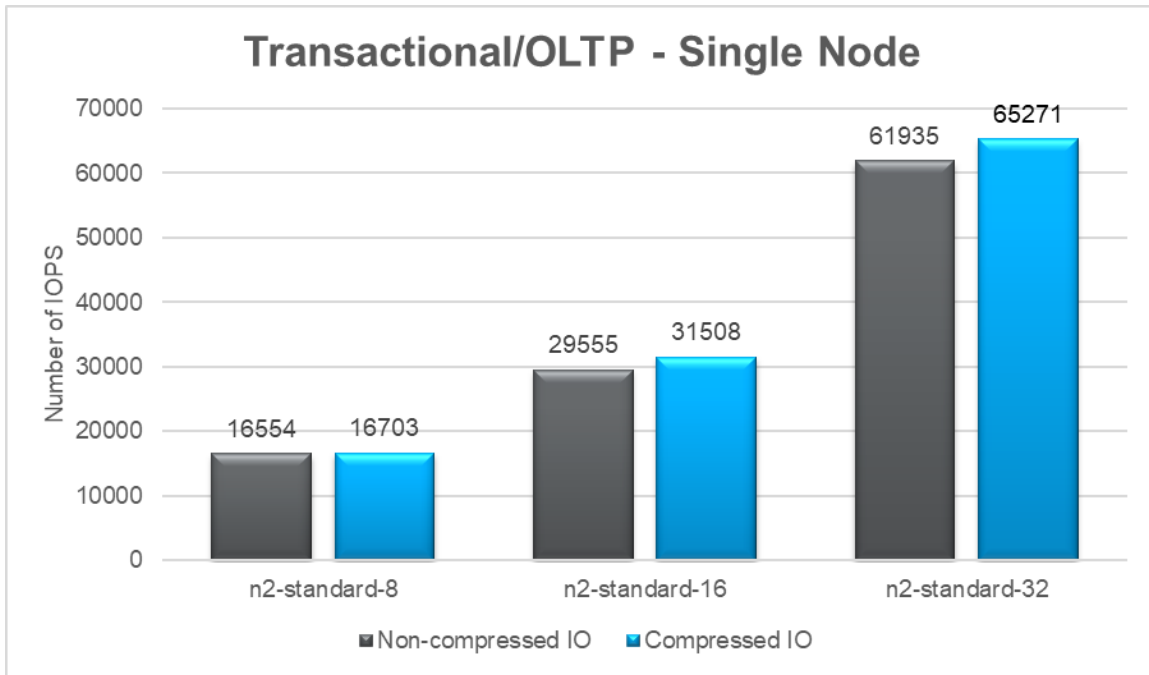


Table 2) Transactional/OLTP workload results for single-node configuration

Machine type (single-node)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	130.49	129.33	16703	16554	3.18	3.22
n2-standard-16	246.16	230.9	31508	29555	1.69	1.8
n2-standard-32	509.93	509.93	65271	61935	1.47	2.06

Figure 3) Transactional/OLTP workload IOPS for a high-availability (HA) configuration.

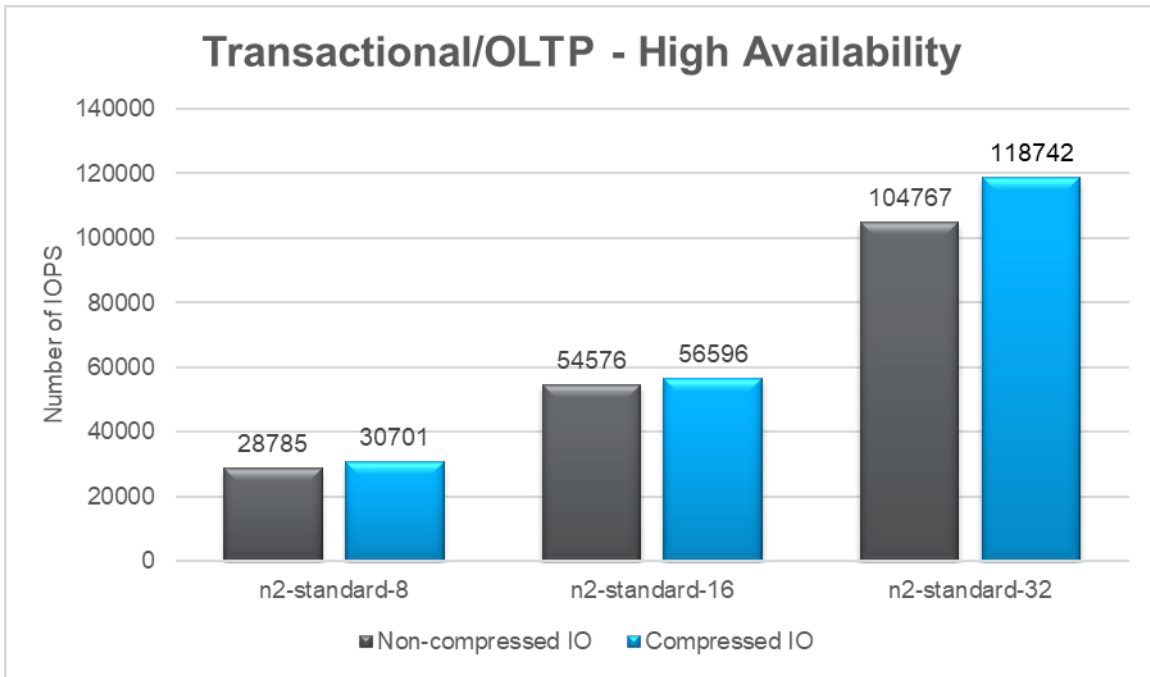


Table 3) Transactional/OLTP workload results for a HA configuration.

Machine type (HA)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	239.85	224.88	30701	28785	3.47	3.7
n2-standard-16	442.16	426.38	56596	54576	1.88	1.95
n2-standard-32	927.67	818.49	118742	104767	1.96	2.54

Streaming reads

The significant metric for this type of workload is throughput, as shown in Figure 4 and Figure 5 and as described in Table 4 and Table 5.

Figure 4) Streaming reads throughput for single-node configuration.




Table 4) Streaming reads workload results for single-node configuration.

Machine types (single node)	Throughput (MBps)	IOPS	Latency (ms)
n2-standard-8	519	8304	4
n2-standard-16	836.56	13385	2.48
n2-standard-32	1293.13	20690	1.61

Figure 5) Streaming reads throughput for HA configuration.

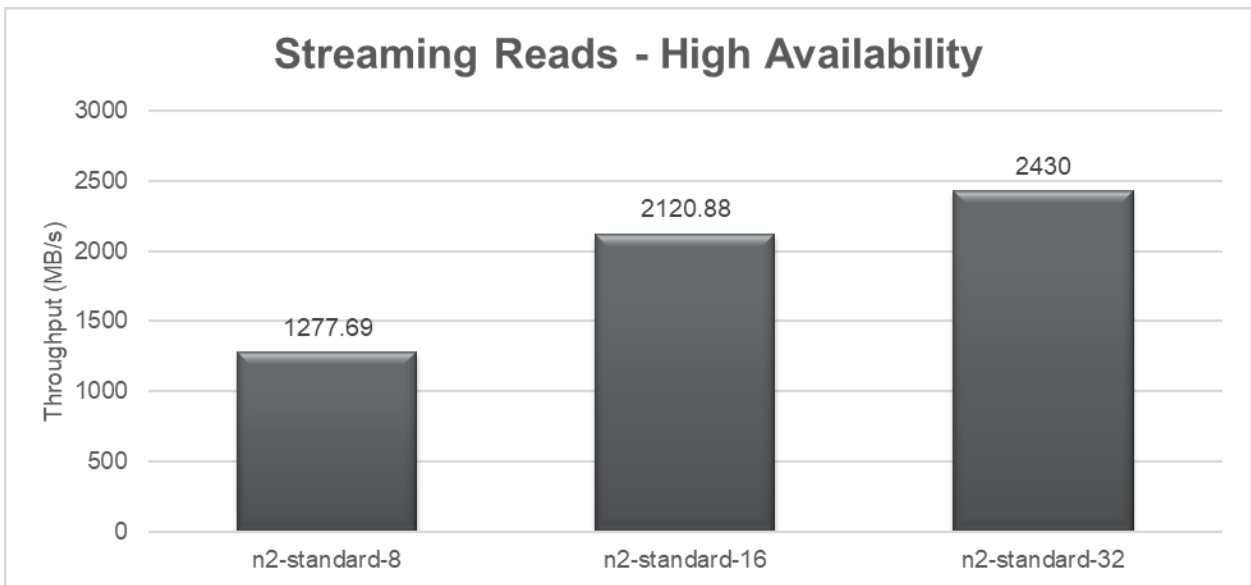


Table 5) Streaming reads workload results for HA configuration.

Machine type (HA)	Throughput (MBps)	IOPS	Latency (ms)
n2-standard-8	1277.69	20443	1.9
n2-standard-16	2120.88	33934	1.18
n2-standard-32	2430	38882	1.69

Streaming writes

The significant metric for this type of workload is throughput, as shown in Figure 6 and Figure 7 and as described in Table 6 and Table 7.

Figure 6) Streaming writes throughput for single-node configuration.

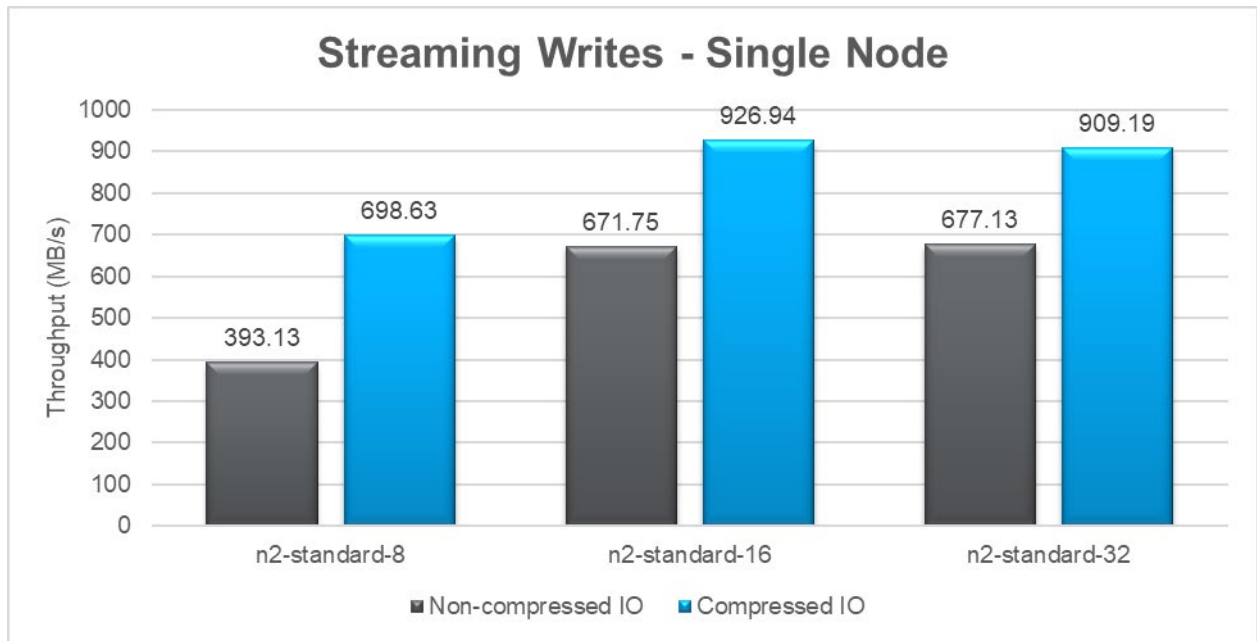


Table 6) Streaming writes workload results for single-node configuration.

Machine Type (single-node)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	698.63	393.13	11178	6290	5.7	5.49
n2-standard-16	926.94	671.75	14831	10748	6.09	7.42
n2-standard-32	909.19	677.13	14547	10834	4.38	7.36

Figure 7) Streaming writes throughput for HA configuration.




Table 7) Streaming writes workload results details for HA configuration.

Machine Type (HA)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	988.56	580.19	15817	9283	8	7.44
n2-standard-16	1249.94	934.75	19999	14956	5.84	8.89
n2-standard-32	1316.94	952.13	21071	15234	6	8.72

Mixed reads/writes

The significant metrics for this type of workload are throughput, as shown in Figure 8 through Figure 11 and as described in Table 8 through Table 11.

Figure 8) Mixed reads/writes 16KB throughput for single-node configuration.




Table 8) Mixed reads/writes 16KB result details for single-node configuration.

Machine type (single-node)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	254.47	230.98	16286	14783	2.45	2.7
n2-standard-16	460.63	387.89	29480	24825	1.76	2.89
n2-standard-32	731.97	594.63	46846	38056	2.55	3.15

Figure 9) Mixed reads/writes 16KB throughput for HA configuration.




Table 9) Mixed reads/writes 16KB result details for HA configuration.

Machine Type (HA)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	406.11	323.95	25991	20733	3.07	3.85
n2-standard-16	789.38	661.95	50520	42365	2.06	3.39
n2-standard-32	1161.44	841.03	74332	53826	3.23	4.45

Figure 10) Mixed reads/writes 32KB throughput for single-node configuration.




Table 10) Mixed reads/writes 32KB result details for single-node configuration.

Machine Type (single-node)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	359.75	278.56	11512	8914	4.51	4.48
n2-standard-16	661.56	518.59	21170	16595	3.39	3.13
n2-standard-32	1151.84	800.56	36859	25618	3.25	3.9

Figure 11) Mixed reads/writes 32KB throughput for HA configuration.




Table 11) Mixed reads/writes 32KB result details for HA configuration.

Machine Type (HA)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	642.25	473.03	20552	15137	5.05	5.27
n2-standard-16	1196.53	914.47	38289	29263	2.5	4.91
n2-standard-32	1788.53	1199.88	57233	38396	4.18	5.2

Random reads

The significant metrics for this type of workload are IOPS, as shown in Figure 12 and Figure 13 and as described in Table 12 and Table 13.

Figure 12) Random reads 8KB IOPS for single-node configuration.




Table 12) Random reads 8KB result details for single-node configuration.

Machine type (single-node)	Throughput (MBps)	IOPS	Latency (ms)
n2-standard-8	116.14	14866	3.53
n2-standard-16	211.73	27101	1.93
n2-standard-32	474.51	60737	1.14

Figure 13) Random reads 8KB IOPS for HA configuration.




Table 13) Random reads 8KB result details for HA configuration.

Machine type (HA)	Throughput (MBps)	IOPS	Latency (ms)
n2-standard-8	231.7	29657	2.16
n2-standard-16	437.18	55959	1.14
n2-standard-32	961.77	123106	1.56

Random writes

The significant metrics for this type of workload are IOPS and latency, as shown in Figure 14 and Figure 15 and as described in Table 14 and Table 15.

Figure 14) Random writes 8KB IOPS for single-node configuration.

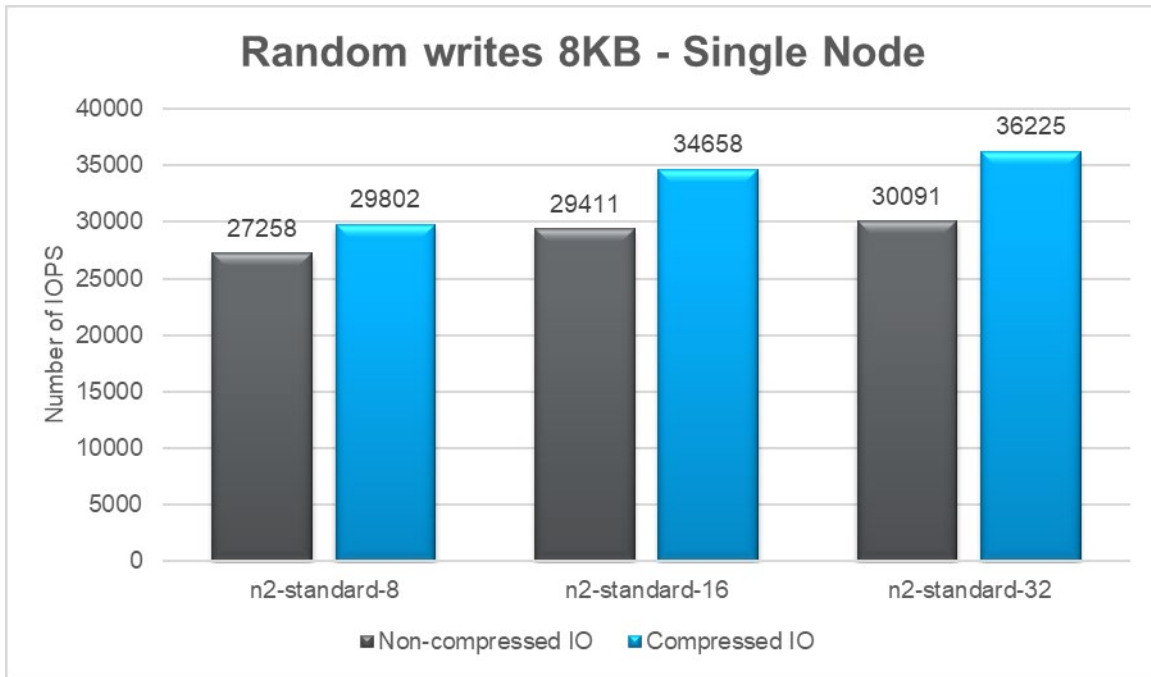


Table 14) Random writes 8KB result details for single-node configuration.

Machine type (single-node)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	232.83	212.95	29802	27258	3.94	4.89
n2-standard-16	270.77	229.77	34658	29411	3.83	4.35
n2-standard-32	283.01	235.09	36225	30091	3.23	4.25

Figure 15) Random writes 8KB IOPS for HA configuration.

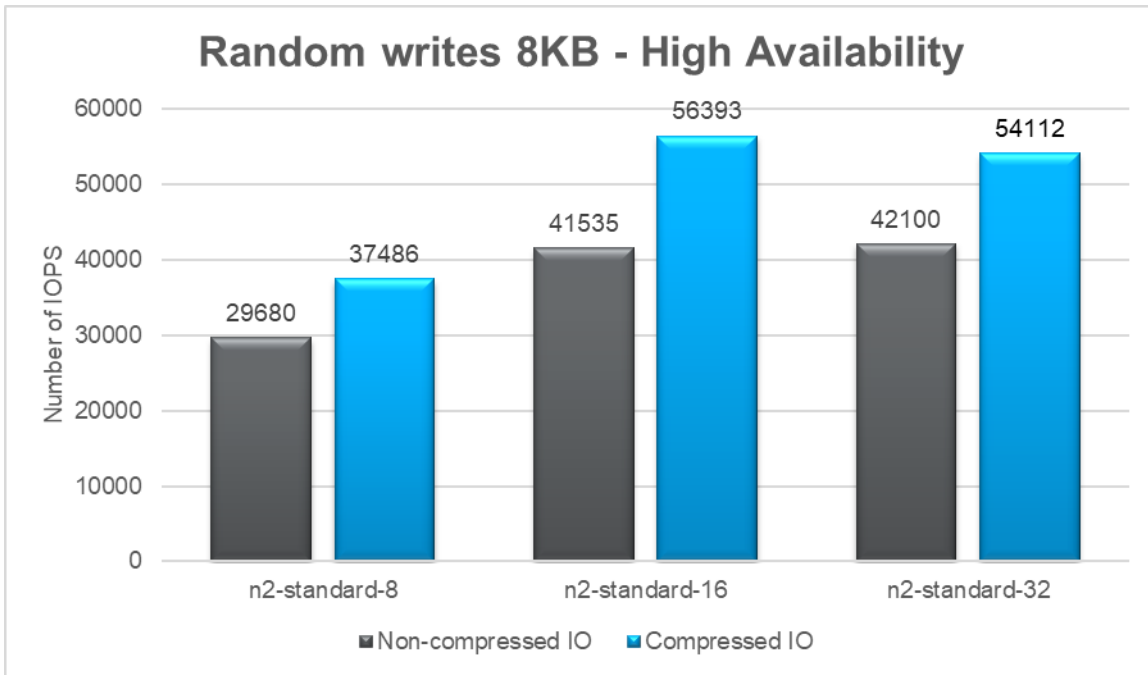


Table 15) Random writes 8KB result details for HA configuration.

Instance type (single-node)	Throughput (MBps)		IOPS		Latency (ms)	
	Yes	No	Yes	No	Yes	No
n2-standard-8	292.86	231.88	37486	29680	5.1	6.46
n2-standard-16	440.57	324.49	56393	41535	4.16	5.64
n2-standard-32	422.75	328.91	54112	42100	4.5	7.5

Acknowledgements

The author gratefully acknowledges the contributions that were made to this technical report by key NetApp team members: Oded Berman, Liron Bar-Mashiah, Aviv Degani, Yaron Haimsohn, Michele Pardini and Ofir Dagan. My sincere appreciation and thanks go to all these individuals, who designed and performed the assessments and provided the insight and expertise that greatly assisted in the creation of this paper.

Where to find additional information

To learn more about the information that is described in this document, review the following documents or websites:

- Cloud Volumes ONTAP (features, architecture, TCO calculator, pricing, and more) <https://cloud.netapp.com/ontap-cloud>
- NetApp BlueXP (formerly NetApp Cloud Manager) and Cloud Volumes ONTAP documentation <https://docs.netapp.com/us-en/cloud-manager-family/>

- Getting Started with Cloud Volumes ONTAP for Google Cloud
https://docs.netapp.com/us-en/occm/task_getting_started_gcp.html
- Fio page on GitHub
<https://github.com/axboe/fio>
- Google Cloud storage option
<https://cloud.google.com/compute/docs/disks/performance>
- Block storage performance page in Google Cloud
<https://cloud.google.com/compute/docs/disks/performance#ssd-pd-performance>
- NetApp BlueXP in Google Cloud Platform Marketplace
<https://console.cloud.google.com/marketplace/details/netapp-cloudmanager/cloud-manager>

Version history

Version	Date	Document version history
Version 1.0	February 2020	Initial release.
Version 1.0.1	January 2023	Updates for ONTAP 9.11.1

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 © 2020 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data—Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

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.

TR-4816-0123