



Technical Report

Performance Characterization of NetApp Cloud Volumes ONTAP for Amazon Web Services

Oded Berman, NetApp
May 2020 | TR-4383

Abstract

This technical report examines the performance of NetApp® Cloud Volumes ONTAP® for Amazon Web Services (AWS) 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

1	Introduction	4
2	Executive Summary.....	4
3	Measuring Storage Performance	4
3.1	Test Methodology	5
3.2	Cloud Volumes ONTAP Configuration	5
3.3	Workloads	7
3.4	Test Environment Diagram.....	8
4	Performance Assessment Results	8
4.1	Single Node.....	9
4.2	High Availability	16
	Acknowledgments	21
	Where to Find Additional Information	22
	Version History	22

LIST OF TABLES

Table 1)	Tested EC2 instance types.....	5
Table 2)	Storage configuration.....	6
Table 3)	Random reads results details.....	9
Table 4)	Transactional/OLTP results details.....	11
Table 5)	Sequential reads results details.....	12
Table 6)	Streaming writes results details.....	14
Table 7)	Mixed reads/writes results details.....	16
Table 8)	Random reads results details.....	17
Table 9)	Transactional/OLTP results details - HA.....	18
Table 10)	Sequential reads results details – HA.....	19
Table 11)	Streaming writes results details - HA.....	20
Table 12)	Mixed reads/writes results details - HA.....	21

LIST OF FIGURES

Figure 1)	Test environment diagram.....	8
Figure 2)	Random reads workload IOPS and latency.....	9
Figure 3)	Transactional/OLTP workload IOPS and latency – Normal Write.....	10
Figure 4)	Transactional/OLTP workload IOPS and latency – High Write vs. Normal Write.....	11
Figure 5)	Sequential reads throughput.....	12
Figure 6)	Sequential writes throughput – Normal Write.....	13

Figure 7) Sequential writes throughput – High Write vs. Normal Write 14
Figure 8) Mixed reads/writes workload throughput and latency – Normal Write 15
Figure 9) Mixed reads/writes workload throughput and latency – High Write vs. Normal Write..... 15
Figure 10) Random reads workload IOPS and latency - HA. 17
Figure 11) Transactional/OLTP workload IOPS and latency - HA. 18
Figure 12) Sequential reads throughput - HA..... 19
Figure 13) Sequential writes throughput – HA.....20
Figure 14) Mixed reads/writes workload throughput and latency – HA.21

1 Introduction

NetApp has a long history of providing leading-edge, high-performance, feature-rich data management solutions through its ONTAP family anywhere across the hybrid multicloud, providing the agility and mobility that organizations need in their digital transformation and cloud journey.

The NetApp® Cloud Volumes ONTAP® software-only storage offering is an enterprise-grade storage solution deployed in the public cloud that runs on top of native cloud compute services and manages native cloud block and object storage services. Today, Cloud Volumes ONTAP can be deployed in any of the “big three” cloud providers: Amazon Web Services (AWS), Azure, and Google Cloud. Cloud Volumes ONTAP optimizes cloud storage costs and performance while enhancing data protection, security, and compliance. It delivers a powerful storage solution for any enterprise workload, such as relational and non-relational databases, DevOps, virtual desktop infrastructure (VDI), Kubernetes, file services, disaster recovery, backup and archiving, and more.

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

This technical report describes the results of performance assessments that were conducted with Cloud Volumes ONTAP for AWS. When deployed in AWS, Cloud Volumes ONTAP uses Amazon Elastic Compute Cloud (EC2), a secure and resizable compute capacity, and Amazon Elastic Block Store (EBS), a high-performance block storage service.

The performance of Cloud Volumes ONTAP is highly dependent on the underlying cloud services being used - the type of EC2 instance, the type and size of EBS volumes, and on the workload. Therefore, several types of workloads, each with different characteristics of I/O operations, were selected to be tested with Cloud Volumes ONTAP using different types of EC2 instances and a non-limiting EBS configuration.

2 Executive Summary

NetApp performed this study to showcase the storage performance and benefits of NetApp Cloud Volumes ONTAP in the AWS cloud. Understanding the performance characteristics of Cloud Volumes ONTAP for AWS is critical to set expectations and to help stakeholders make educated decisions about which of their workloads can be handled by Cloud Volumes ONTAP and to understand what the required cloud resources are.

The results are presented in section 4, Performance Assessment Results. The results provide some basic understanding of the Cloud Volumes ONTAP performance capabilities in the AWS cloud and its dependency on the underlying infrastructure and the services that are used.

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

3 Measuring Storage Performance

NetApp performed these assessments to measure the performance of Cloud Volumes ONTAP 9.7 in single-node and high-availability (HA) configurations with 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.

3.1 Test Methodology

For this study, NetApp used Flexible I/O tester (FIO) version 3.16, 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, read the [FIO's documentation](#).

To get the most out of Cloud Volumes ONTAP, multiple FIO clients were used in the various assessments. The I/O workloads were generated by multiple `m5.2xlarge` EC2 instances (each with eight virtual CPUs [vCPUs] and 32GB of memory) running Red Hat Enterprise Linux 8.

All tests were performed with NFS v4, 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 a NetApp FlexVol® volume, an allocation of storage space that serves as a data container, hosted on Cloud Volumes ONTAP. 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, multiple FlexVol volumes were configured and accessed by the clients. 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 EC2 instances that NetApp used in the study, including the instances that Cloud Volumes ONTAP used, were created in the EU-WEST-1 (Ireland) region, in the same availability zone, Virtual Private Cloud, and subnet.

3.2 Cloud Volumes ONTAP Configuration

For this study, Cloud Volumes ONTAP was configured in both single node and high availability (HA) pair as an active-active configuration. Since the performance of Cloud Volumes ONTAP depends greatly on the underlying cloud services that are used (EC2 and EBS), different EC2 types were used and a non-limiting EBS configuration. The EC2 types not only differ in their specifications, but they also incur different costs by AWS and by NetApp and belong to different pricing options (PayGo and BYOL) and PayGo licensing packages. The complete list of supported instances can be found in the [supported configurations for Cloud Volumes ONTAP 9.7 in AWS](#).

The specific combination of resources determines the overall performance of an EC2 instance. For Cloud Volumes ONTAP, the amount of vCPU effects the general performance, whereas the memory size directly affects the read performance. The EBS bandwidth acts as a throughput throttling factor that is independent of read and write performance levels. All of the supported instance types for Cloud Volumes ONTAP are EBS optimized and have dedicated bandwidth to Amazon EBS, with different maximal throughput, depending on the instance type. For additional information on the instance types, visit the [Amazon instance types](#) page and [Amazon EBS-optimized instances](#) in the EC2 user guide.

EC2 Instances

The following table (Table 1) displays the EC2 instance types used by Cloud Volumes ONTAP that were tested:

Table 1) Tested EC2 instance types.

Model	vCPU		Memory (GiB)	Storage Type (GiB)	Network Bandwidth (Gbps)	EBS Bandwidth (Mbps)
	EC2	ONTAP ³				
c5n.18xlarge ⁴	72	48	192	EBS-only	100	19,000

Model	vCPU		Memory (GiB)	Storage Type (GiB)	Network Bandwidth (Gbps)	EBS Bandwidth (Mbps)
	EC2	ONTAP ³				
c5n.9xlarge	36	36	96	EBS-only	50	9,500
c5.9xlarge	36	36	72	EBS-only	10	9,500
c5d.18xlarge ²	72	48	144	2 x 900 NVMe SSD	25	19,000
c5d.9xlarge ²	36	36	72	1 x 900 NVMe SSD	10	9,500
r5dn.12xlarge ^{2,4}	48	48	384	2 x 900 NVMe SSD	10	9,500
r5.12xlarge	48	48	384	EBS-only	10	9,500
r5.8xlarge	32	24	256	EBS-only	10	6,800
r5.2xlarge ¹	8	8	64	EBS-only	10	4,750
m5.16xlarge	64	48	256	EBS-only	20	13,600
m5d.12xlarge ²	48	48	192	2 x 900 NVMe SSD	10	9,500
m5d.8xlarge ²	32	24	128	2 x 600 NVMe SSD	10	6,800
m5.4xlarge	16	16	64	EBS-only	10	4,750
m5.2xlarge ¹	8	8	32	EBS-only	10	4,750

Note:

1. These instance types can support maximum performance for 30 minutes at least once every 24 hours. If you have a workload that requires sustained maximum performance for longer than 30 minutes, select a different instance type.
2. These instance types, in addition to EBS, include local NVMe storage, which Cloud Volumes ONTAP uses as NetApp Flash Cache. NetApp Flash Cache speeds access to data through real-time intelligent caching. It is effective for random read-intensive workloads, including databases, email, and file services.
3. Currently, Cloud Manager deploys Cloud Volumes ONTAP with 16 vCPUs as a maximum limit, however additional vCPUs are supported. Changing the number of vCPUs currently requires a manual procedure.
4. Instances are not officially supported yet.

Storage Configuration

All Cloud Volumes ONTAP nodes were configured with multiple EBS General Purpose SSD (gp2) volumes as shown in the following table (Table 2) – For more information, see diagram in section 3.4, Test Environment Diagram.

Table 2) Storage configuration.

Resources	Single Node	High Availability Node
# of GP2 EBS volumes	12 (each 6TB in size)	6 (each 6TB in size)
# of aggregates	2 (each with 6 EBS volumes)	1 (with 6 EBS volumes)
# of FlexVol volumes in aggregate	2	2

Using General Purpose SSD is the recommended best practice because it balances price performance for a wide variety of workloads and transactional data. Cloud Volumes ONTAP supports other types of EBS volumes, such as EBS Provisioned IOPS SSD (io1), Throughput Optimized HDD (st1) and Cold HDD (sc1). For additional information about Amazon EBS, visit [Amazon EBS volume types](#).

Write Speed

Cloud Volumes ONTAP's [Write Speed](#) option, dramatically affects write performance and can be set to either **Normal** (default) or **High**.

When the write speed is set to `Normal`, data is first written to persistent 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 storage. 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.

All instances were tested with Write Speed set to `Normal`. To demonstrate the Write Speed effect on write performance, two single node Cloud Volumes ONTAP instances were tested with the option set to `High`. In section 4 Performance Assessment Results, the values for both instances, in both Normal and High write speeds are presented in the applicable workloads.

3.3 Workloads

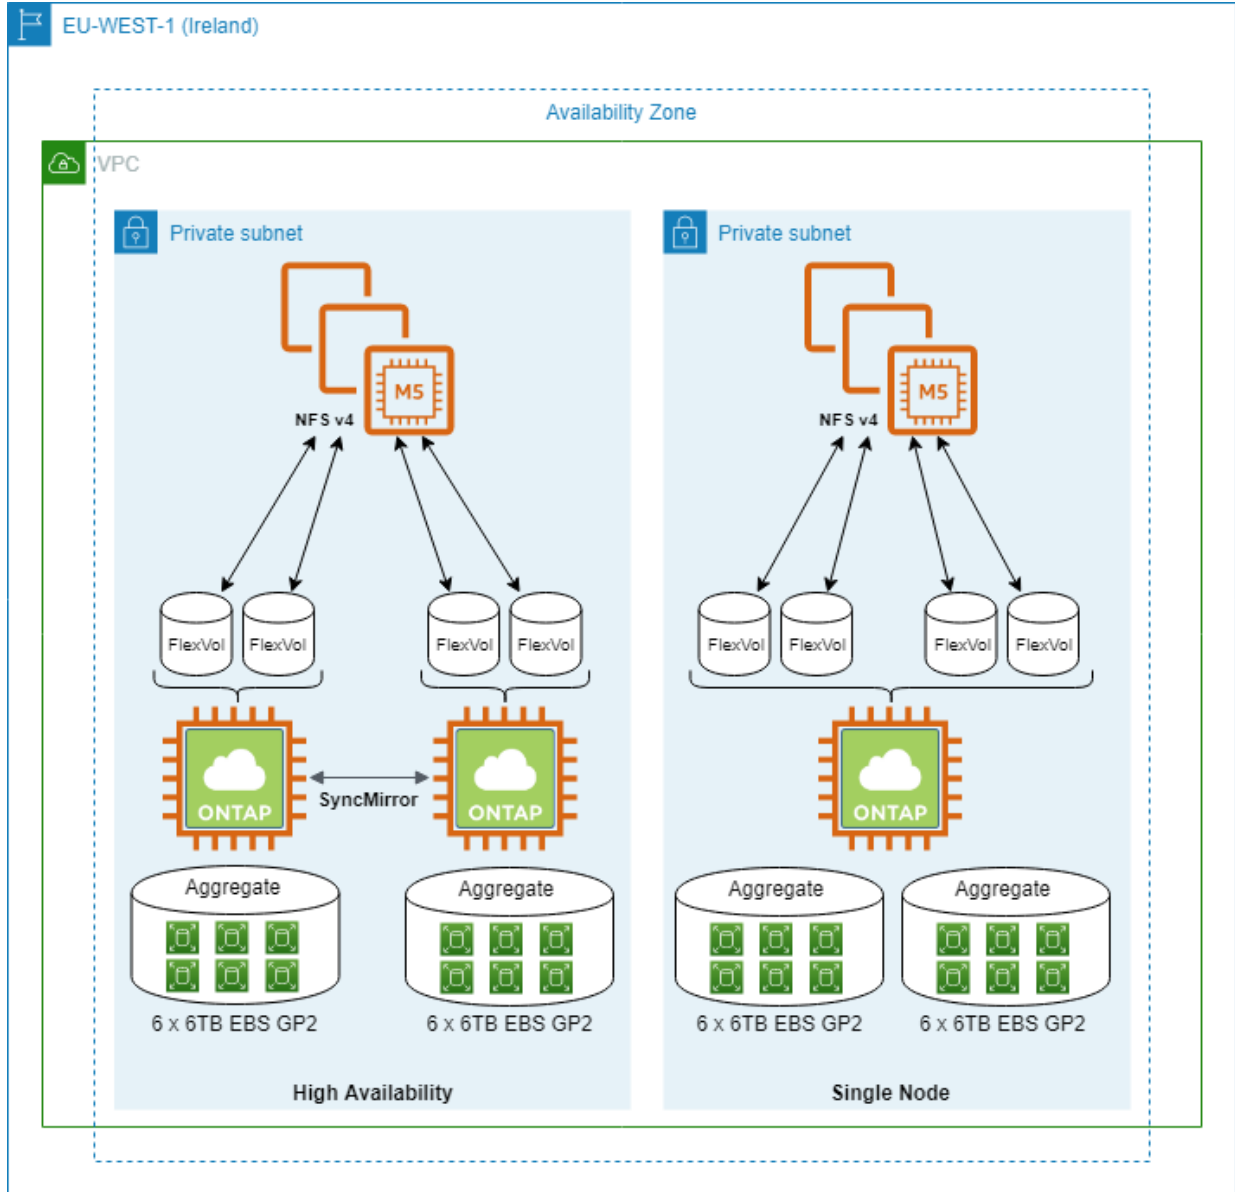
In this study, five different workloads were configured and tested. Each tested workload consisted of a unique I/O mixture and access pattern that simulated workloads that are used in widely deployed applications:

- **Random reads** (4K block size 100% reads, and 100% random access I/O). A random read is a disk access pattern whereby small (4K) blocks of data are read from random locations. This workload is used primarily within the context of benchmarking how effective a device is at quickly retrieving small pieces of data from random locations.
- **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.
- **Sequential reads** (64KB block size, 100% reads, and 100% sequential access I/O). In a streaming reads workload, only sequential 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.
- **Sequential writes** (64KB block size, 100% writes, and 100% sequential access I/O). In a streaming writes workload, only sequential 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** (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.

3.4 Test Environment Diagram

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

Figure 1) Test environment diagram.



4 Performance Assessment Results

This section describes the values of the performance metrics that were obtained in the process of assessing NetApp Cloud Volumes ONTAP for AWS for a single-node and HA pair as an active-active configuration deployed in a single availability zone.

4.1 Single Node

The following figures and tables describe the performance metrics of Cloud Volumes ONTAP for AWS in a single node configuration in applicable workloads.

Random Reads

The significant metrics for this type of workload are IOPS and latency, as shown in the following figure and table (Figure 2 and Table 3).

Figure 2) Random reads workload IOPS and latency.




Table 3) Random reads results details.

Model	IOPS	Throughput (MBps)	Latency (ms)
m5.2xlarge	N/A	N/A	N/A
c5d.18xlarge	200263	801	1.52
m5d.12xlarge	132312	529	3.29
c5d.9xlarge	130848	523	3.3
r5dn.12xlarge	125855	503	3.55
c5.18xlarge	99280	397	3.56
c5n.18xlarge	95282	381	4.84
m5.16xlarge	76933	308	6.03
r5.12xlarge	72920	292	6.08
m5d.8xlarge	70663	283	6.19
m5.4xlarge	63331	253	7.35

Model	IOPS	Throughput (MBps)	Latency (ms)
c5n.9xlarge	51753	207	9.09
r5.8xlarge	50819	203	9.22
c5.9xlarge	50537	202	9.29

Transactional/OLTP

The significant metrics for this type of workload are IOPS and latency, as shown in the following figures and tables (Figure 3 and Figure 4 and Table 4).

Figure 3) Transactional/OLTP workload IOPS and latency – Normal Write.




Figure 4) Transactional/OLTP workload IOPS and latency – High Write vs. Normal Write.




Table 4) Transactional/OLTP results details.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5d.18xlarge	96236	770	4.7
r5dn.12xlarge ¹	94729 (192188)	758 (1537)	5.9 (2)
m5d.12xlarge	91566	733	2.5
c5n.18xlarge ¹	90002 (123134)	720 (985)	5.3 (3.8)
c5.18xlarge	88948	712	4.29
m5.16xlarge	88388	707	2.6
m5d.8xlarge	85164	681	5.5
c5d.9xlarge	82776	662	5.41
r5.12xlarge	79544	636	5.6
r5.2xlarge	61799	494	7.39
m5.4xlarge	61235	490	7.53
c5n.9xlarge	55253	442	8.49
c5.9xlarge	54899	439	5.45
r5.8xlarge	44995	360	10.33

1. High write speed values in parenthesis.

Sequential Reads

The significant metric for this type of workload is throughput, as shown in the following figure and table (Figure 5 and Table 5).

Figure 5) Sequential reads throughput.

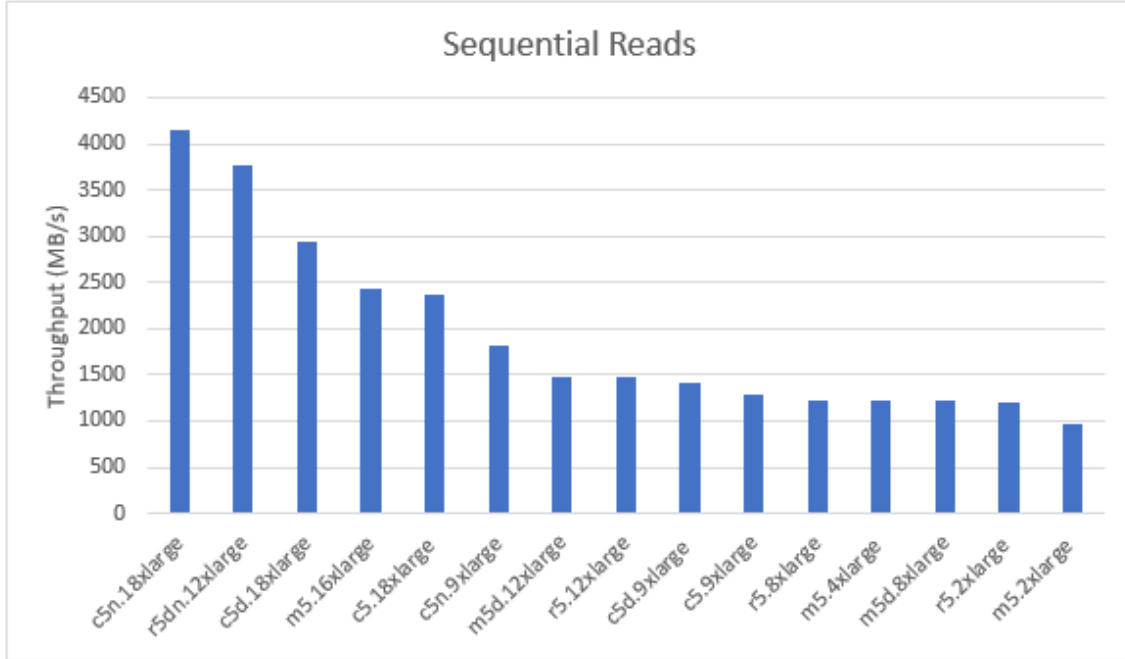


Table 5) Sequential reads results details.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5n.18xlarge	64949	4157	5
r5dn.12xlarge	58799	3763	6.9
c5d.18xlarge	45857	2935	6
m5.16xlarge	38139	2441	2
c5.18xlarge	36949	2365	7.3
c5n.9xlarge	28451	1821	16.1
m5d.12xlarge	22961	1470	0.46
r5.12xlarge	22941	1468	0.88
c5d.9xlarge	22082	1413	2.69
c5.9xlarge	19954	1277	11.2
r5.8xlarge	19104	1223	2
m5.4xlarge	19026	1218	4.19
m5d.8xlarge	19023	1217	2.07

Model	IOPS	Throughput (MBps)	Latency (ms)
r5.2xlarge	18866	1207	1.36
m5.2xlarge	15214	974	0.6

Sequential Writes

The significant metric for this type of workload is throughput, as shown in the following figures and table (Figure 6 and Figure 7 and Table 6).

Figure 6) Sequential writes throughput – Normal Write.

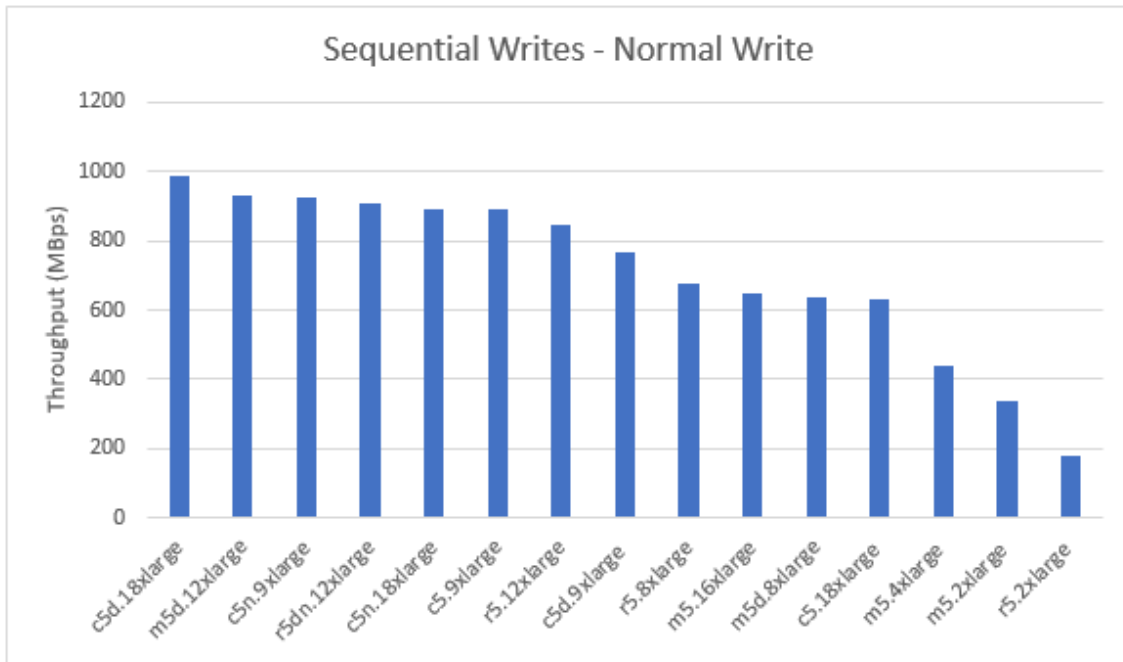


Figure 7) Sequential writes throughput – High Write vs. Normal Write.

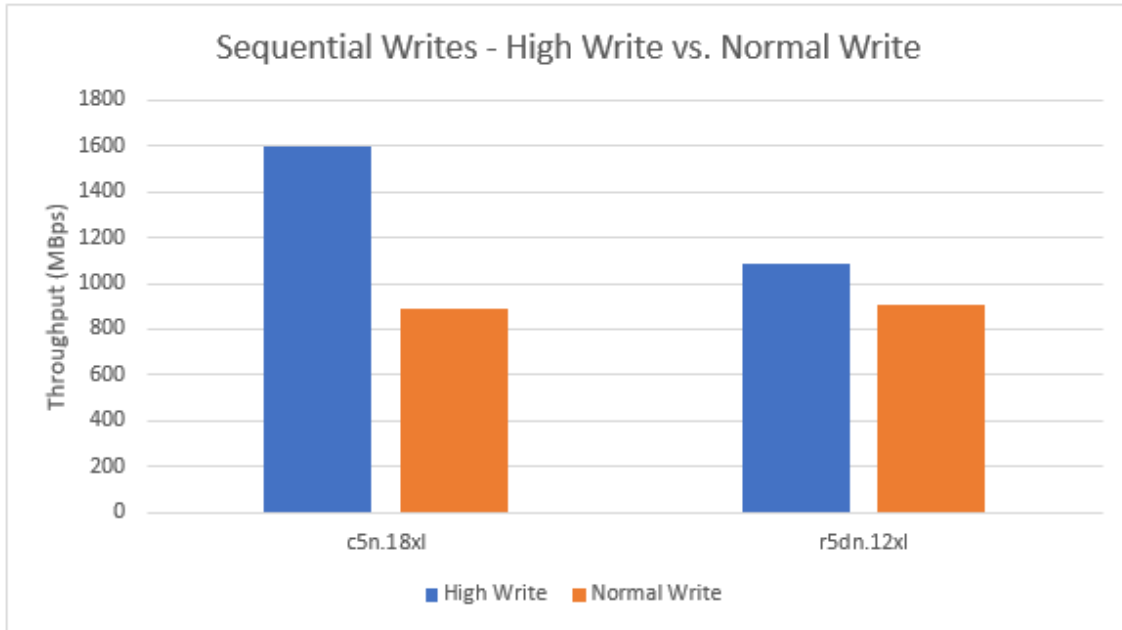


Table 6) Streaming writes results details.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5d.18xlarge	15465	990	32
m5d.12xlarge	14530	930	32
c5n.9xlarge	14496	928	33
r5dn.12xlarge ¹	14183 (16989)	908 (1087)	47 (11)
c5n.18xlarge ¹	13948 (24951)	893 (1597)	35 (8)
c5.9xlarge	13911	890	34
r5.12xlarge	13198	845	36
c5d.9xlarge	12013	769	9.23
r5.8xlarge	10557	676	40.53
m5.16xlarge	10138	649	46.29
m5d.8xlarge	9987	639	43.06
c5.18xlarge	9882	632	8.83
m5.4xlarge	6843	438	37.42
m5.2xlarge	5264	337	14.5
r5.2xlarge	2745	176	44.53

1. High write speed values in parenthesis.

Mixed Reads/Writes

The significant metrics for this type of workload are throughput and latency, as shown in the following figures and table (Figure 8 and Figure 9 and Table 7).

Figure 8) Mixed reads/writes workload throughput and latency – Normal Write.

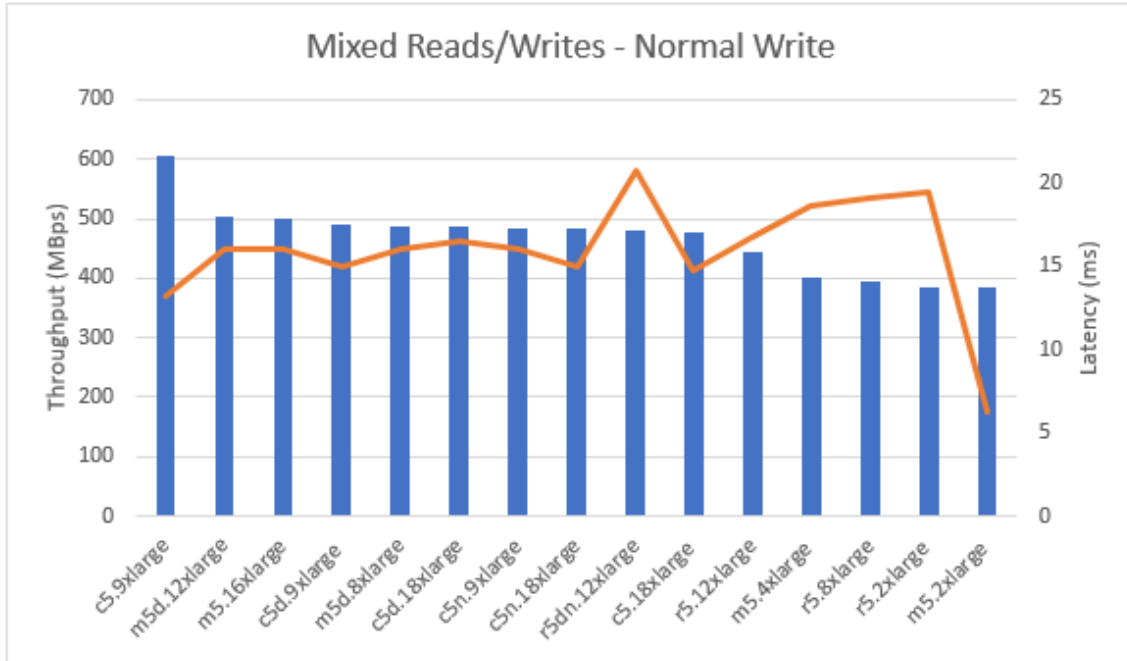


Figure 9) Mixed reads/writes workload throughput and latency – High Write vs. Normal Write.

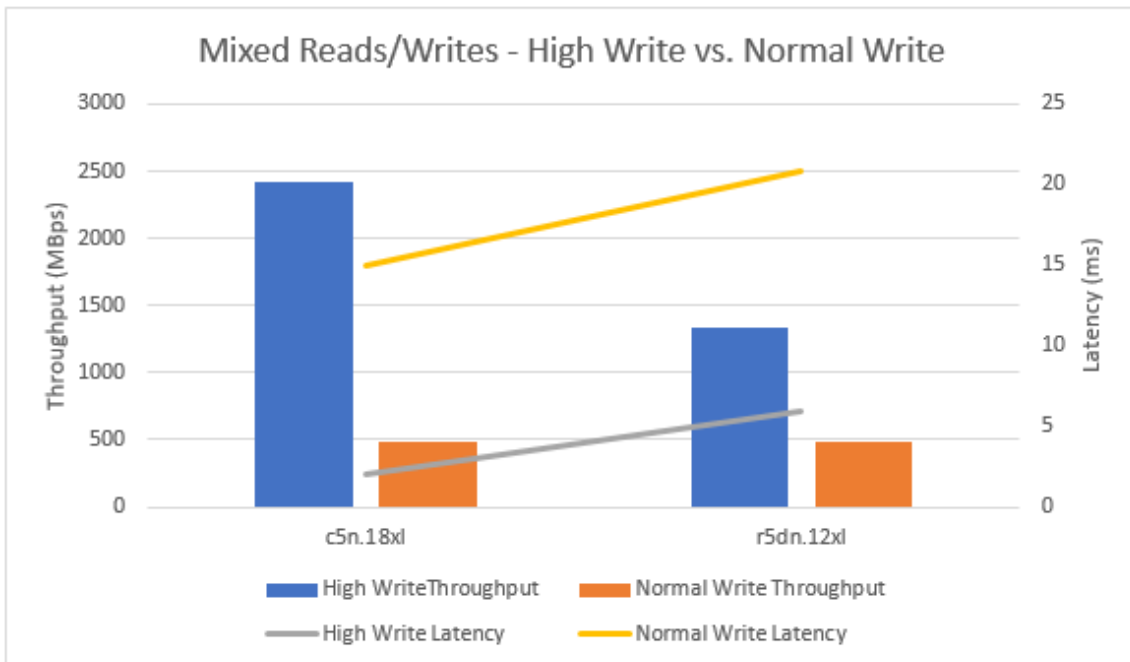


Table 7) Mixed reads/writes results details.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5.9xlarge	37788	605	13.2
m5d.12xlarge	31403	502	16
m5.16xlarge	31274	500	16
c5d.9xlarge	30618	490	14.98
m5d.8xlarge	30449	487	16
c5d.18xlarge	30371	486	16.5
c5n.9xlarge	30295	485	16
c5n.18xlarge ¹	30265 (151271)	484 (2420)	15 (2)
r5dn.12xlarge ¹	30085 (83415)	481 (1334)	20.8 (5.9)
c5.18xlarge	29812	477	14.72
r5.12xlarge	27714	443	16.72
m5.4xlarge	25103	402	18.65
r5.8xlarge	24556	393	19.09
r5.2xlarge	24059	385	19.45
m5.2xlarge	23991	384	6.2

1. High write speed values in parenthesis.

4.2 High Availability

The following figures and tables describe the performance metrics of Cloud Volumes ONTAP for AWS in a high availability (HA) pair as an active-active configuration deployed in a single availability zone. Results displayed are the sum of both nodes in the HA pair.

Random Reads

The significant metrics for this type of workload are IOPS and latency, as shown in the following figure and table (Figure 10 and Table 8).

Figure 10) Random reads workload IOPS and latency - HA.




Table 8) Random reads results details.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5d.18xlarge	542012	2168	0.29
m5d.12xlarge	502212	2009	1.02
m5d.8xlarge	367894	1472	1.85
r5dn.12xlarge	353734	1415	2.26
r5.12xlarge	352144	1409	2.57
m5.16xlarge	309754	1239	2.83
c5n.18xlarge	267360	1069	3.1
c5n.9xlarge	217368	869	4.21
r5.8xlarge	216202	865	4.12
c5.9xlarge	210914	844	4.33

4.2.2 Transactional/OLTP

The significant metrics for this type of workload are IOPS and latency, as shown in the following figure and table (Figure 11 and Table 9).

Figure 11) Transactional/OLTP workload IOPS and latency - HA.

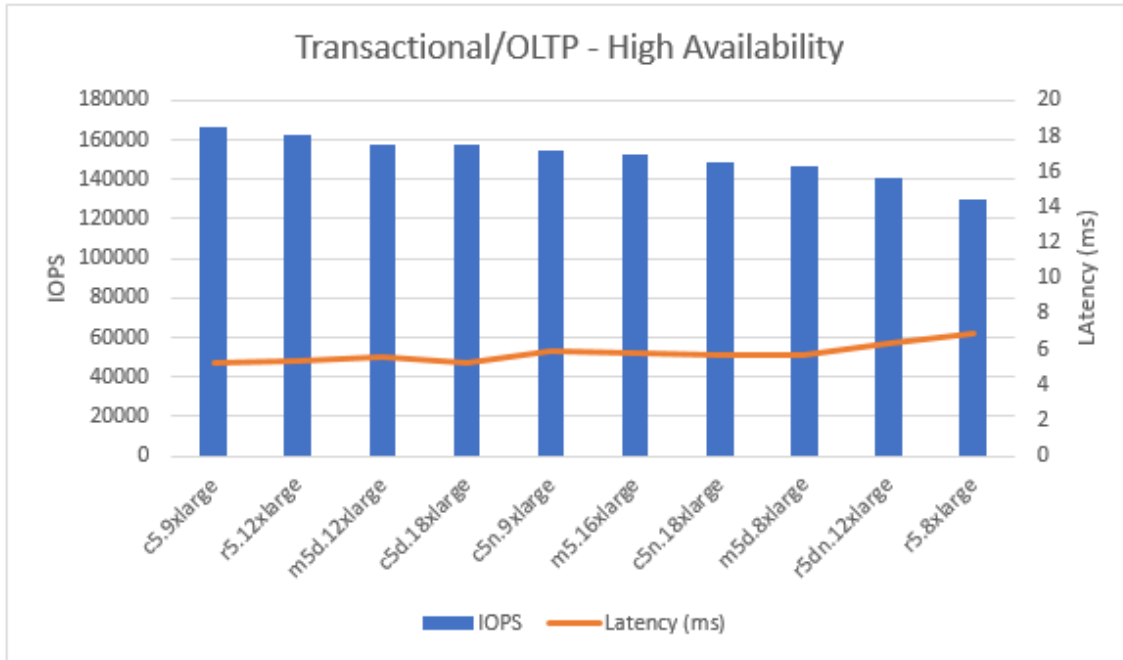


Table 9) Transactional/OLTP results details - HA.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5.9xlarge	166352	1331	5.25
r5.12xlarge	162506	1300	5.37
m5d.12xlarge	157636	1261	5.54
c5d.18xlarge	157466	1260	5.25
c5n.9xlarge	154934	1239	5.85
m5.16xlarge	152418	1219	5.78
c5n.18xlarge	148458	1188	5.68
m5d.8xlarge	146626	1173	5.72
r5dn.12xlarge	141016	1128	6.39
r5.8xlarge	130118	1041	6.86

Sequential Reads

The significant metric for this type of workload is throughput, as shown in the following figure and table (Figure 12 and Table 10).

Figure 12) Sequential reads throughput - HA.

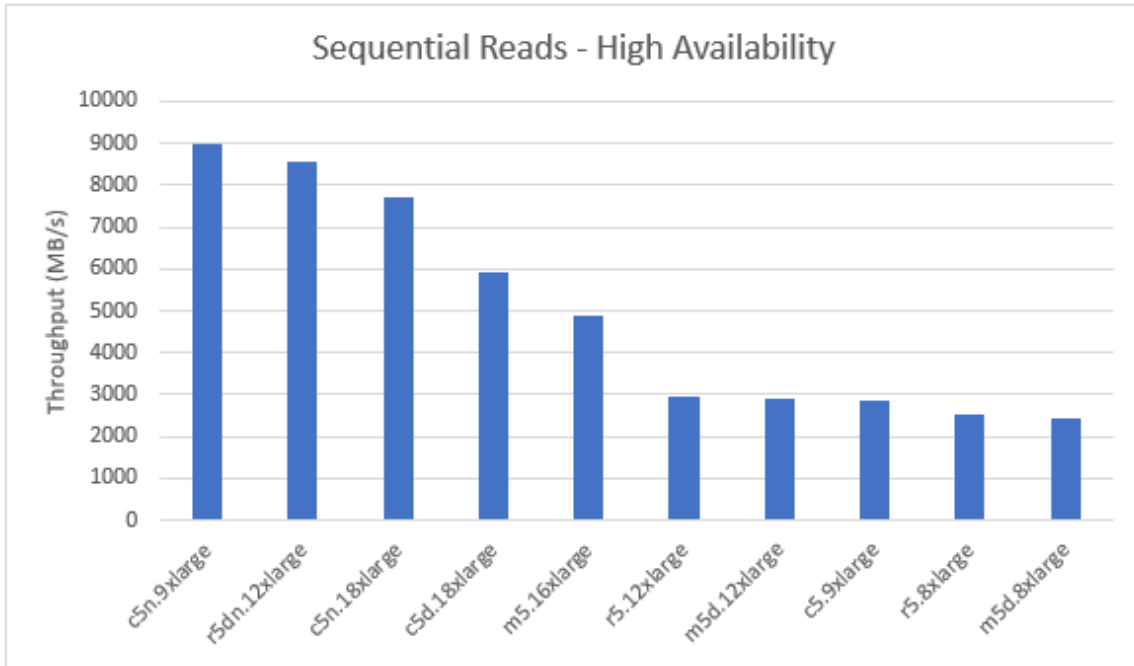


Table 10) Sequential reads results details – HA.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5n.9xlarge	140198	8973	2.81
r5dn.12xlarge	133472	8542	4.68
c5n.18xlarge	120786	7730	3.9
c5d.18xlarge	92392	5913	2.74
m5.16xlarge	76334	4885	2.04
r5.12xlarge	45828	2933	0.11
m5d.12xlarge	45726	2926	0.41
c5.9xlarge	44526	2850	0.73
r5.8xlarge	39594	2534	0.18
m5d.8xlarge	38272	2449	0.71

Sequential Writes

The significant metric for this type of workload is throughput, as shown in the following figure and table (Figure 13 and Table 11).

Figure 13) Sequential writes throughput – HA.

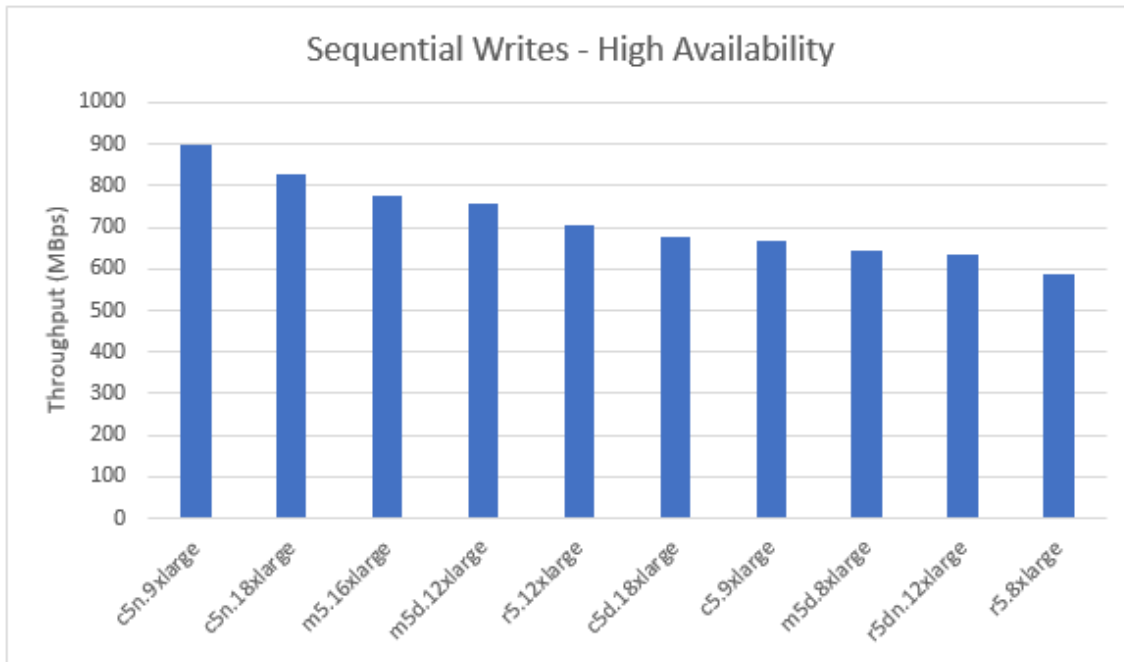


Table 11) Streaming writes results details - HA.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5n.9xlarge	14010	897	37.64
c5n.18xlarge	12904	826	41.3
m5.16xlarge	12140	777	44.56
m5d.12xlarge	11792	755	72.17
r5.12xlarge	11010	705	47.39
c5d.18xlarge	10556	676	53.84
c5.9xlarge	10402	666	49.92
m5d.8xlarge	10050	643	50.36
r5dn.12xlarge	9922	635	53.84
r5.8xlarge	9146	585	91.58

Mixed Reads/Writes

The significant metrics for this type of workload are throughput and latency, as shown in the following figure and table (Figure 14 and Table 12).

Figure 14) Mixed reads/writes workload throughput and latency – HA.



Table 12) Mixed reads/writes results details - HA.

Model	IOPS	Throughput (MBps)	Latency (ms)
c5.9xlarge	54320	869	16.56
r5.12xlarge	53134	850	17.46
m5d.12xlarge	51146	818	17.99
c5n.18xlarge	49790	797	17.72
c5n.9xlarge	48120	770	19.53
m5d.8xlarge	47890	766	19.24
m5.16xlarge	47294	757	19.4
c5d.18xlarge	45450	727	16.79
r5dn.12xlarge	42122	674	19.15
r5.8xlarge	41418	663	22.37

Acknowledgments

The author gratefully acknowledges the contributions that were made to this technical report by key NetApp team members: Ofir Dagan, Ishai Avramovich, Oryan Rachum, Aviv Degani, Michele Pardini, and Yuval Kalderon. 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 Cloud Manager and Cloud Volumes ONTAP documentation
<https://docs.netapp.com/us-en/occm/>
- Getting Started with Cloud Volumes ONTAP for AWS
https://docs.netapp.com/us-en/occm/task_getting_started_aws.html
- Enabling Flash Cache on Cloud Volumes ONTAP in AWS
https://docs.netapp.com/us-en/occm/task_enabling_flash_cache.html#whats-flash-cache
- Planning your Cloud Volumes ONTAP write speed
https://docs.netapp.com/us-en/occm/task_planning_your_config.html#choosing-a-write-speed
- Latest FIO documentation
<https://fio.readthedocs.io/en/latest/>
- Amazon EBS features
<https://aws.amazon.com/ebs/features/>
- Amazon EC2 Instance Types
<https://aws.amazon.com/ec2/instance-types/>
- Cloud Manager – Deploy and Manage NetApp Cloud Data Services in AWS Marketplace
<https://aws.amazon.com/marketplace/pp/B07QX2QLXX>

Version History

Version	Date	Document Version History
Version 3	April 2020	Revision of performance with Cloud Volumes ONTAP 9.7 for AWS
Version 2.3	Jan 2019	Updated performance with Cloud Volumes ONTAP 9.5
Version 2.2	July 2018	Updated performance with Cloud Volumes ONTAP 9.4 Updated instance types Added Tiering performance
Version 2.1	February 2018	Updated performance with Cloud Volumes ONTAP 9.3
Version 2.0	June 2017	Updated performance with Cloud Volumes ONTAP 9.2 Added performance of high availability
Version 1.2	May 2017	Additional tests
Version 1.1	January 2016	Updated with tests of additional instance types
Version 1.0	February 2015	Initial release

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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. 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).

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-4383-0520