



Technical Report

# Performance Characterization of NetApp Cloud Volumes ONTAP in Amazon Web Services with Application Workloads

NetApp Data Fabric Group, NetApp  
October 2018 | TR-4383

## **Abstract**

This technical report examines the performance and fit of application workloads to NetApp® Cloud Volumes ONTAP® instances that are running in Amazon Web Services (AWS) EC2.

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Differences Among Instance Types</b>	<b>3</b>
<b>3</b>	<b>Test Configurations</b>	<b>4</b>
<b>4</b>	<b>Summary of the Test Results</b>	<b>4</b>
4.1	OLTP workload	5
4.2	Streaming read workload	5
4.3	Streaming write workload	5
4.4	Analytics workloads	6
4.5	Data on FabricPool	6
<b>5</b>	<b>Conclusion</b>	<b>6</b>
	<b>Where to Find Additional Information</b>	<b>7</b>
	<b>Version History</b>	<b>7</b>

## LIST OF TABLES

Table 1)	Capabilities of supported instance types	3
Table 2)	OLTP workload IOPS and latency	5
Table 3)	Streaming read workload throughput	5
Table 4)	Streaming write workload throughput	6
Table 5)	Analytics workload throughput	6

# 1 Introduction

To help NetApp customers select the most appropriate solutions for their IT infrastructure, NetApp provides up-to-date documentation describing its products. This technical report describes the results of performance tests for the NetApp Cloud Volumes ONTAP data management software running on an Amazon Web Services (AWS) EC2 instance. NetApp partners, customers, and employees can use this information to make informed decisions about which workloads are appropriate for Cloud Volumes ONTAP.

The test configuration environments described in this report consist of the following components:

- Cloud Volumes ONTAP clusters, both single node and high availability (HA) pairs, of the following instance types:
  - c4.8xlarge (Premium license)
  - c4.4xlarge (Premium license)
  - m4.2xlarge (Standard license)
  - m4.xlarge (Explore license)
- A Microsoft Windows Server 2012 R2 host running Iometer. For each test, the instance type used for the Iometer machine matched the Cloud Volumes ONTAP system instance type.
- The iSCSI protocol for block workloads.

## 2 Differences Among Instance Types

The [Amazon instance types](#) page describes the capabilities of the different instance types in terms of network, CPU, and memory.

The specific combination of these components is what determines the overall performance of an instance. For Cloud Volumes ONTAP, the CPU capability contributes to general performance, whereas memory contributes heavily to read performance. The network capability acts as a throughput limiting factor that is independent of read and write performance levels. All the supported instance types for Cloud Volumes ONTAP are Elastic Block Store (EBS) optimized and thus deliver dedicated bandwidth to Amazon EBS, with different maximal throughput, depending on the instance type.

Table 1) Capabilities of supported instance types.

Instance	CPU	RAM	Network	EBS Max Throughput for 128KB Streaming Read Workload (MBps)
m4.xlarge	4	16GB	High	93.75
m4.2xlarge	8	32GB	High	125
m4.4xlarge	16	64GB	High	250
r4.xlarge	4	30.5GB	Up to 10 Gigabit	109
r4.2xlarge	8	61GB	Up to 10 Gigabit	218
c4.4xlarge	16	30GB	High	250
c4.8xlarge	36	60GB	10 Gigabit	500

### 3 Test Configurations

All test configurations used the iSCSI protocol for block I/O connectivity. The tests focused on the following:

- Iometer was used to generate I/O workloads. We used different number of LUNs and different number of outstanding I/O operations to see the effect of concurrency level on IOPS, throughput, and latency.
- Normal writing speed was used for single-node systems. Under normal write speed, data is written directly to disk, minimizing the likelihood of data loss if an unplanned system outage occurs. Conversely, when high write speed is used, data is buffered in memory before it is written to disk. This configuration provides faster write performance but increases the potential for data loss if an unplanned system outage occurs. High write speed is recommended if fast write performance is required and you can tolerate the risk of data loss due to an unplanned system outage. For example, data loss might be handled by the application.
- We tested single availability zone HA Cloud Volumes ONTAP systems and single-node deployments. In the HA tests, a single Iometer client generated load against multiple LUNs that all resided on one of the nodes in the cluster, rather than having LUNs being split between the nodes with two clients generating load against them. Read performance can be doubled when LUNs are divided between nodes.
- We tested different instance types. The instance types tested not only differed by their hardware characteristics but also incurred different costs by AWS and by NetApp (they belong to different licensing packages). An instance type from each licensing package was picked for the tests.
- All the tests were performed with gp2 EBS volumes. Cloud Volumes ONTAP supports additional EBS volume types: st1, sc1, and io1. Although gp2 and io1 volumes provide the best performance results, gp2 volumes are the recommended configuration because they are much more cost-effective.
- Tests were performed on a freshly deployed system. Each workload was run for 40 minutes without prior warmup.
- Cloud Volumes ONTAP was tested on four different instance types:
  - **m4.xlarge** – included in the Explore license that allows up to 2TB of storage
  - **m4.2xlarge** – included in the Standard license that allows up to 10TB of storage
  - **c4.4xlarge** – included in the Premium license that allows up to 368TB of storage
  - **c4.8xlarge** – included in the Premium license. This instance type has strong hardware capabilities. Therefore, its AWS pricing is the highest of all supported instance types.

### 4 Summary of the Test Results

Each tested configuration consisted of a unique workload that is representative of the workloads used in widely deployed POSIX applications:

- **OLTP workload.** 8KB block size, 100% random access I/O, and a mixture of 70% reads. This workload simulates database applications (SAP, Oracle, SQL) and OLTP servers.
- **Streaming reads.** 64KB block size, 100% reads, and 100% sequential access I/O. This workload simulates applications such as media servers (for example, video on demand) and virtual tape libraries.
- **Streaming writes.** 64KB block size, 100% writes, 100% sequential access I/O. This workload simulates applications such as media capture, virtual type libraries, medical imaging, archiving, backup, and video surveillance.
- **Analytics.** 16KB block size, 50% reads, 100% random access I/O. Analytics workloads are ad hoc in nature. They involve both read and write operations and require high throughput and low latency.

The test workloads differed markedly. Collectively, however, they represent workloads that can be considered for an Cloud Volumes ONTAP deployment. Because the workloads were so different, the results of each test are discussed separately in this section.

## 4.1 OLTP Workload

The 8KB block size workload simulates an OLTP transactional database. Transactional workloads tend to be read-heavy as data about an item is retrieved, but they involve a smaller number of writes as transactions are committed. Transactional workloads are highly sensitive to write latency, especially for writes to the transaction log. Typically, it is best for log writes not to exceed 10ms of latency. Lower latency is always better. Table 2 shows the results of 8KB blocks, 70% reads, 100% random access I/O for the different instance types.

Table 2) OLTP workload IOPS and latency.

Instance Type	IOPS for Single Node (Ops/s)	Latency for Single Node (msec)	IOPS for HA pair (Ops/s)	Latency for HA pair (msec)	License
c4.8xlarge	39216	2	25486	3	Premium
c4.4xlarge	43877	2	26753	4	Premium
m4.2xlarge	20745	2	10757	6	Standard
m4.xlarge	9863	13	6727	10	Explore

## 4.2 Streaming Read Workload

We tested workload with large contiguous read request size that simulates applications such as Media Servers (for example, video on demand) and virtual tape libraries. The workload consisted of 64KB block size, 100% reads, 100% sequential access I/O.

The most significant measurement for this type of workload is throughput, as Table 3 shows.

Table 3) Streaming read workload throughput.

Instance Type	Throughput for Single Node (MBps)	Throughput for HA pair (MBps)	License
c4.8xlarge	350	340	Premium
c4.4xlarge	360	330	Premium
m4.2xlarge	115	115	Standard
m4.xlarge	86	86	Explore

Cloud Volumes ONTAP has robust performance for streaming reads. The differences between the instance types are mainly caused by the networking capabilities of the instance and the number of CPUs available. For the HA pair tests, all LUNs were resident on the same node for the tested configuration. Splitting the LUNs between the two nodes can increase the performance even further, with the potential to double it.

## 4.3 Streaming Write Workload

We tested a workload with large contiguous write request size that simulates applications such as media capture, virtual tape libraries, medical imaging, archiving, backup, video surveillance and reference data. The workload consisted of 64KB block size, 100% writes, 100% sequential access I/O.

The most significant measurement for this type of workload is throughput, as Table 4 shows.

Table 4) Streaming write workload throughput.

Instance Type	Throughput for Single Node (MBps)	Throughput for HA pair (MBps)	License
c4.8xlarge	470	220	Premium
c4.4xlarge	215	204	Premium
m4.2xlarge	110	106	Standard
m4.xlarge	43	34	Explore

Cloud Volumes ONTAP performed well for the streaming write workload. Here too, single node performed better than HA pair, due to the propagation of each write to the other node

#### 4.4 Analytics Workloads

We tested a workload with analytics characteristics. Analytics consist of read and write mixture, for which throughput and latency are the most important measurements. The workload we tested consisted of 16KB block size, 50% read, 100% random access I/O.

Table 5) Analytics workload throughput.

Instance Type	Throughput for Single Node (MBps)	Latency for Single Node (msec)	Throughput for HA pair (MBps)	Latency for HA pair (msec)	License
c4.8xlarge	370	3	195	6	Premium
c4.4xlarge	320	3	240	6	Premium
m4.2xlarge	169	5	86	10	Standard
m4.xlarge	54	25	48	38	Explore

The results show that Cloud Volumes ONTAP is a good fit for analytics workloads, because it can maintain high throughput under strict latency constraints.

#### 4.5 Data on FabricPool

Beginning with ONTAP 9.4, all inactive or cold data in the active file system can be tiered to the AWS S3. We tested workloads on the volumes which had the relevant tiering policy (auto, backup), to cool inactive data to the AWS S3.

##### Writing to AWS S3

We worked with iometer and the iSCSI protocol to stream a write workload on a volume with tiering capacity set to 'auto' while the data was written to EBS. In the background, data that was considered cold was transferred to S3. We did not see any change in write performance.

##### Reading from AWS S3

We worked with the FIO tool on Linux as a client and the NFS protocol, reading from a mounted volume, while all the inactive data was cooled to the S3. Running 64K BS, random read, average read size was 16K. We saw degradation of about 30% in throughput and IOPS, compared to reading all data from the EBS disks.

## 5 Conclusion

Cloud Volumes ONTAP was found as good fit for OLTP workloads, streaming read workloads, streaming write workloads and analytics workloads. NetApp has a long history of providing high-performance and feature-rich storage systems. Cloud Volumes ONTAP extends this legacy to AWS. With Cloud Volumes

ONTAP, NetApp continues to develop leading-edge storage solutions that provide the agility and mobility that current NetApp customers need and that future NetApp customers will want. Cloud Volumes ONTAP is part of a family of products that stretch from the private cloud to the hybrid cloud to the public cloud and that run the NetApp ONTAP data management software. Understanding the performance characteristics of Cloud Volumes ONTAP is critical for setting our customers' expectations and enabling their continued success.

## Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- NetApp Product Documentation  
<https://docs.netapp.com>

## Version History

Version	Date	Document Version History
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 © 2018 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.