



Technical Report

# NetApp AFF8080A EX Storage Efficiency and Performance with Microsoft SQL Server 2014

Ron Pratt, Scott Lane, NetApp  
June 2016 | TR-4516

## Abstract

This report provides a storage efficiency and performance summary for NetApp® and partner systems engineers who are interested in assessing SQL Server database storage efficiency and performance with a NetApp AFF8080A EX storage system.

NetApp All Flash FAS (AFF) systems uniquely combine the extreme performance capability of flash media with the industry-leading NetApp Data ONTAP® software to provide performance acceleration, operational agility, best-in-class data protection, and business continuance for database deployments.

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	ONTAP FlashEssentials Empowers All Flash FAS Performance	3
1.2	NetApp ONTAP 9	3
1.3	Storage Efficiency	4
1.4	Microsoft SQL Server 2014	5
<b>2</b>	<b>Executive Summary</b>	<b>5</b>
<b>3</b>	<b>Measuring Storage Efficiency and Performance</b>	<b>6</b>
3.1	Test Methodology	6
3.2	Hardware and Software	6
3.3	Network Design	8
3.4	Database Layout and Storage Provisioning Design	8
3.5	Workload Design	9
3.6	Storage Efficiency Test Results	10
3.7	Performance Test Results	11
<b>4</b>	<b>Best Practices</b>	<b>12</b>
<b>5</b>	<b>Conclusion</b>	<b>12</b>
	<b>References</b>	<b>13</b>

## LIST OF TABLES

Table 1)	SQL Server hardware and software components	6
Table 2)	NetApp storage system hardware and software	7
Table 3)	Database layout	9

## LIST OF FIGURES

Figure 1)	Network design	8
Figure 2)	ONTAP 9 storage efficiency	11
Figure 3)	AFF8080A EX SQL Server 2014 performance	12

# 1 Introduction

This document describes the storage efficiency and performance of the NetApp AFF8080A EX storage system with Microsoft SQL Server 2014 workloads.

## 1.1 ONTAP FlashEssentials Empowers All Flash FAS Performance

NetApp ONTAP FlashEssentials is the power behind the performance and efficiency of All Flash FAS (AFF). ONTAP is a well-known operating system, but what is not widely known is that Data ONTAP, with its WAFL® (Write Anywhere File Layout) file system, is natively optimized for flash media.

ONTAP FlashEssentials encapsulates key features that optimize solid-state-drive (SSD) performance and endurance, including the following:

- NetApp data-reduction technologies, including inline compression, inline deduplication, and inline data compaction, can provide significant space savings. Savings can be further increased by using NetApp Snapshot® and NetApp FlexClone® technologies. Studies based on customer deployment have shown that total data-reduction technologies have enabled up to 933 times space savings.
- Inline data compaction provides continued innovation beyond compression and deduplication, further increasing storage efficiency.
- Coalesced writes to free blocks maximize performance and flash media longevity.
- Flash-specific read-path optimizations provide consistent low latency.
- Parallelized processing handles more requests at once.
- Software-defined access to flash maximizes deployment flexibility.
- New advanced drive partitioning (ADP) increases storage efficiency and further increases usable capacity by almost 20%.
- Data-fabric readiness enables live workload migration between flash and hard-disk-drive tiers, on premise or to the cloud.
- Quality-of-service capability safeguards service-level objectives in multiworkload and multitenant environments.

## 1.2 NetApp ONTAP 9

NetApp ONTAP 9 is a major advance in the industry's leading enterprise data management software. This software can integrate the best of next-generation and traditional technologies, incorporating flash, cloud, and software-defined architectures while building a foundation for the data fabric. Plus, new customers and existing Data ONTAP 8.3 environments can quickly and easily use the rich data services delivered by ONTAP 9.

An essential feature for SQL Server databases deployed on shared enterprise storage is the capability to deliver consistent and dependable high performance. High performance must be coupled with nondisruptive operations, high availability, scalability, and storage efficiency. Customers can depend on ONTAP 9 and AFF to provide these essential elements.

Built on the Data ONTAP unified scale-out architecture, AFF consistently meets or exceeds the high-performance demands of SQL Server databases. AFF also provides rich data management capabilities, such as integrated data protection and nondisruptive upgrades and data migration. These features allow customers to eliminate performance silos and seamlessly integrate AFF into a shared infrastructure. ONTAP 9 delivers enhanced inline deduplication and completely new inline data compaction capability that significantly reduces the amount of flash storage required, with no impact on system performance. ONTAP 9 also provides industry-leading ecosystem integration with database applications that makes administration of databases and storage systems far more efficient when compared with other flash storage solutions on the market.

## 1.3 Storage Efficiency

Simply stated, storage efficiency enables you to store the maximum amount of data within the smallest possible space at the lowest possible cost. The following NetApp storage efficiency technologies can help you realize maximum space savings:

- **Inline compression.** Data compression reduces the disk space required, regardless of storage protocol, application, or storage tier. Inline compression also reduces the data that must be moved to SSDs, thereby reducing the wear on SSDs. Furthermore, the reduction in the amount of data being written can deliver an increase in overall performance.
- **Inline and always-on deduplication.** Data deduplication cuts storage requirements by reducing redundancies in primary, backup, and archival data. Inline deduplication of zeros speeds up VM provisioning by 20% to 30%. Improvements to inline deduplication in ONTAP 9 provide additional efficiency by extending elimination of duplicate data to blocks in memory and SSDs.<sup>1</sup>
- **Inline data compaction.** NetApp inline data compaction in ONTAP 9 provides significant storage savings by compressing and coalescing small I/O into single block writes. Doing so further reduces the disk space required and associated wear on SSDs.
- **Snapshot technology.** NetApp Snapshot technology provides low-cost, instantaneous, point-in-time space-efficient copies of the file system (volume) or LUN by preserving ONTAP architecture and WAFL consistency points without affecting performance. NetApp SnapManager® integrates with SQL Server Virtual Device Interface (VDI) to create application-consistent Snapshot copies of production-level SQL Server databases with no downtime for the production database.
- **Thin provisioning.** Thin provisioning, implemented by NetApp at the NetApp FlexVol® volume and LUN level, defers storage purchases by keeping a common pool of free storage available to all applications.
- **Thin replication.** Thin replication is at the center of the NetApp data protection software portfolio, which includes NetApp SnapMirror® and NetApp SnapVault® software. SnapVault thin replication enables more frequent backups that use less storage capacity because no redundant data is moved or stored. SnapMirror thin replication protects business-critical data while minimizing storage capacity requirements.
- **RAID DP.** NetApp RAID DP® technology protects against double disk failure without sacrificing performance or adding disk-mirroring overhead.
- **FlexClone volumes.** FlexClone virtual cloning reduces the need for storage by enabling multiple, instant, space-efficient, writable copies.
- **Advanced drive partitioning.** Advanced SSD partitioning with ONTAP 9 increases usable capacity by almost 20%.<sup>2</sup>

The total effective capacity is the total amount of logical data that can be stored based on savings from compression, deduplication, data compaction, Snapshot copies, and FlexClone copies. The efficiency ratio is calculated by taking the sum of all of the logical data used in each aggregate and comparing it to the sum of physical blocks used to support the logical data. The higher the ratio, the greater the space savings.

---

<sup>1</sup> <https://www.netapp.com/us/media/tr-4476.pdf>

<sup>2</sup> <https://www.netapp.com/us/media/ds-3582.pdf>

## 1.4 Microsoft SQL Server 2014

The NetApp solution for Microsoft SQL Server delivers the performance, efficiency, manageability, scalability, and data protection that IT organizations need to meet even the most stringent SQL Server business requirements. NetApp helps organizations to maximize the full potential of their SQL Server investment by accelerating database application performance and speeding application deployment from months to weeks. Companies that deploy NetApp solutions for SQL Server also realize the availability and flexibility benefits required to be competitive in today's business environment.

## 2 Executive Summary

To showcase the benefits of the AFF8080A EX running ONTAP 9, NetApp conducted the following studies:

- **Measuring ONTAP 9 storage efficiency.** NetApp measured the storage efficiency of several customer production databases after data reduction by inline compression, inline deduplication, and inline data compaction.
- **Measuring AFF8080A EX Microsoft SQL Server 2014 performance.** NetApp measured the data throughput, input/output operations per second (IOPS), and average latency of the AFF8080A EX storage controllers running ONTAP 9 with an industry-standard SQL Server 2014 online transaction processing (OLTP) workload. All inline storage efficiency features were enabled.

NetApp conducted two studies to measure the storage efficiency and performance of an AFF8080A EX storage system running ONTAP 9 under an industry-standard SQL Server 2014 OLTP workload. There were two goals. The first was to determine the storage savings on several samples of customer production data from the storage efficiencies delivered by ONTAP 9. The second goal was to determine the peak IOPS and average storage latency of the AFF8080A EX storage system while running the SQL Server workload.

We measured the overall storage efficiency savings of a customer production SQL Server database deployed on ONTAP 9. In our test configuration, a total space savings of 15:1 was observed. This total savings is attributed to the overall efficiency effectiveness of Snapshot copies, inline compression, and inline data compaction.

Additionally, we ran an OLTP Transaction Processing Performance Council Benchmark E (TPC-E) workload on a 2-node AFF8080A EX cluster with 2 DS2246 shelves containing a total of 48 800GB SSD drives. We tested our cluster at a range of load points that drove the storage to peak CPU utilization. At each load point, we collected information about the storage IOPS and latency.

The SQL Server performance test demonstrated that the cluster IOPS increased from 4K IOPS to a peak performance of 350K IOPS. For all load points up to 339K IOPS, the test configuration maintained consistent storage latencies at or below 1ms.

## 3 Measuring Storage Efficiency and Performance

NetApp conducted two studies to measure the storage efficiency and performance of an AFF8080A EX storage system running ONTAP 9. The following sections describe the methodology and design considerations used to test the AFF8080A EX while running a standard SQL Server workload.

### 3.1 Test Methodology

In the storage-efficiency study, we measured the efficiency of both the data generated by the TPC-E workload generator as well as a customer production database.

To measure the storage efficiency of data generated by the TPC-E workload generator, we used TPC-E to populate a SQL Server database. We created one aggregate on each of the storage nodes of the AFF8080A EX. After installing the client operating systems and SQL Server database software, and configuring all software, we used TPC-E to populate the SQL Server database as it would prior to a TPC-E performance test run. After the database population was completed, we measured the storage efficiency of the workload files using the `storage aggregate show-efficiency ONTAP` command on the cluster.

For our storage-efficiency testing, we also tested a 1245GB customer production database. To measure the storage efficiency of this database, we first created two aggregates, one on each storage node. Each aggregate contained one volume. We copied the database into a volume on one node and then used the Linux command `dd` to copy the data into the volume on the second node. Using the `dd` command, we used a block size of 8KB to simulate the native SQL Server IO block size. Copying this dataset allowed the ONTAP inline efficiencies to operate on the data. We used the ONTAP `storage aggregate show-efficiency` command to collect our storage-efficiency measurements.

For our performance study, we used the industry-standard TPC-E OLTP warehouse transaction workload against the SQL Server test configuration. The workload generated a read-write ratio of approximately 90:10 against the SQL Server databases in the test configuration.

We created a SQL Server environment with eight database servers connected through Fibre Channel (FC) to the AFF8080A EX. Each server created a 1.4TB SQL Server database on the AFF8080A EX storage system for a total data size of 11.2TB across the eight SQL Server databases. Using the eight SQL Server databases and an OLTP load generator, we measured the peak performance of the storage system by generating a workload designed to maximize the storage system utilization. We then reran the test while ramping down the server count from eight. This allowed us to gather performance metrics at a range of different load points.

### 3.2 Hardware and Software

For these studies, we configured eight SQL Server 2014 database servers on eight Fujitsu RX300s7 servers. We connected the eight servers to a two-node AFF8080A EX, with four servers on each node of the storage system, using 8Gb FC on the server side and 16Gb FC on the storage side. Each AFF node was connected to one DS2246 disk shelf with 800GB SSD drives following NetApp cabling best practices.

Table 1 and Table 2 list the hardware and software components used for the SQL Server performance test configuration.

Table 1) SQL Server hardware and software components.

Hardware/Software Components	Details
SQL Server 2014 servers	8 Fujitsu RX300s7
Server operating system	Microsoft Windows 2012 R2 Standard Edition
SQL Server database version	Microsoft SQL Server 2014 Enterprise Edition

Hardware/Software Components	Details
Processors per server	2 6-core Xeon E5-2630 @ 2.30GHz
Physical memory per server	128GB
FC network	8Gb FC with multipathing
FC host bus adapter (HBA)	QLogic QLE2562 dual-port PCIe FC HBA
Dedicated public 1GbE ports for cluster management	2 Intel 1350GbE ports
8Gb FC switch	Brocade 6510 24-port
10GbE switch	Cisco Nexus 5596


Table 2) NetApp storage system hardware and software.

Hardware/Software Components	Details
Storage system	AFF8080A EX configured as a high availability (HA) active-active pair
ONTAP version	9.0
Total number of drives	48
Drive size	800GB
Drive type	SSD
FC target ports	8 16Gb ports (4 per node)
Ethernet ports	4 10Gb ports (2 per node)
Storage virtual machines (SVMs)	1 across both node aggregates
Ethernet logical interfaces (LIFs)	4 1Gb management LIFs (2 per node connected to separate private VLANs)
FC LIFs	8 16Gb data LIFs

### 3.3 Network Design

This section provides the network connectivity details for the tested configurations. The network diagram in Figure 1 shows that the FC SAN was deployed with a Brocade 6510 16Gb FC switch. Each server and storage controller node had four ports connected into the SAN. The multiple ports used in the FC SAN configurations provided both HA through multipathing and increased bandwidth. At no point in the testing did the network connectivity create a bottleneck.

Figure 1) Network design.



### 3.4 Database Layout and Storage Provisioning Design

Table 3 summarizes the layout for the SQL Server database.

In the performance test, we used eight 1.4TB databases to host the simulated warehouse transaction environment. Each storage system node housed a single aggregate containing 23 800GB SSD drives that were subdivided into RAID DP groups, plus one spare drive. We used the first controller's data

aggregate to store the first four database files, logs, and temp files. The second controller's data aggregate was used to store the other four database files, logs, and temp files.

As a best practice, NetApp recommends splitting each database across multiple volumes with each volume assigned to a unique LUN.

For the AFF8080A EX, we used ONTAP ADP to create two partitions on each shelf: one for the root aggregate and one for the data aggregate.

**Table 3) Database layout.**

Storage	Aggregate Name	Volume Name	LUN Size (GB)	Vol Size (GB)	File Size (GB)	Description
Per controller						Used ADP
	aggr0	root		55		Total aggregate size = 55GB
	aggr_sql					21 data + 2 parity RAID DP + 1 spare Total aggregate size = 12.7TB
Per DB		Db1	200	300	168	Database file 1
		Db2	200	300	168	Database file 2
		Db3	200	300	168	Database file 3
		Db4	200	300	168	Database file 4
		Db5	200	300	168	Database file 5
		Db6	200	300	168	Database file 6
		Db7	200	300	168	Database file 7
		Db8	200	300	168	Database file 8
		Temp1			1	Temp files
		Log1	300	300	26	Log files

The database layout shown in Table 3 was repeated for each of the eight SQL Server databases. For each SQL Server database, the data files were contained in a separate LUN within a separate volume, for a total of 72 volumes and LUNs across all 8 databases.

Igroups containing all of the FC initiators for each server was configured on the cluster. The LUNS were formatted by using Disk Management with GPT and the NTFS file system using a 64KB block size. The FC SAN was configured on the Brocade switch without zoning. ONTAP provided asymmetric logical unit assignment (ALUA) communication to the initiators so that optimal paths were used for host I/O access according to the multipathing input/output (MPIO) load-balancing policies on the host.

### 3.5 Workload Design

The OLTP workload that we used in testing was generated by a Microsoft partner toolkit that created a TPC-E workload. The toolkit, based on the BenchCraft TPC-E toolkit, was used to simulate an OLTP workload of a brokerage firm. Each database server applied the workload to a full set of SQL Server databases, logs, and temp files. We analyzed the workload pattern and found it to be random and composed of 90% reads and 10% writes with an average block size of 8K. NetApp recommends setting the SQL Server tunings to meet your specific performance requirements.

To establish the performance boundaries, we tested the environment up to a storage system saturation point. When we reached peak storage utilization, we reduced the workload and measured the IOPS and latency at each load point. At this load level, we verified that the storage system and SQL Server instances could maintain steady-state behavior without failure. We also made sure that there were no bottlenecks across servers or networking systems.

After the peak load using eight database workloads was achieved, we noted the total IOPS, throughput, and storage and server latencies. We then generated a load of 75% peak by using six of the eight available SQL Server instances, a load of 50% peak by using four of the eight available SQL Server instances, a load of 25% peak by using two of the eight available SQL Server instances, and a load of 12.5% peak by using one of the eight available SQL Server instances. We tested at each load level and provided the results in this report.

**Note:** We took care in these test steps to simulate real database and customer workloads, but we acknowledge that workloads vary across databases. In addition, these test results were obtained in a closed lab environment with no competing workloads on the same infrastructure. In a typical shared-storage infrastructure, other workloads share resources. Your results might vary from those found in this report.

### 3.6 Storage Efficiency Test Results

We measured the overall storage efficiency savings of two SQL Server databases on ONTAP 9. The two databases that we used were:

- 5.6TB of data generated from the TPC-E workload generator
- 1245GB database files from a customer production environment


For the TPC-E workload generator storage efficiency testing, we used two disk shelves, each containing 24 800GB SSD drives, for a total raw storage of 38.4TB, as mentioned in section 3.2 “Hardware and Software”. Of the 24 drives per shelf, 23 drives were used to create one single RAID DP aggregate and 1 was left as a spare drive. We used ONTAP ADP to partition the drives, install the root partition, and create the data partition by using RAID DP to provide data redundancy, which resulted in total usable storage of 25.5TB. Of this usable space, we used the TPC-E workload generator to create 5.6TB of data for the SQL Server database. The amount of compression that can be achieved depends highly on the specific contents of the data that is written and stored in the database.

As shown in Figure 2, from the 5.6TB of TPC-E storage, we measured a space savings of 1.9TB from inline compression. This resulted in a storage efficiency ratio of 1.5:1 from inline compression. In this dataset, we found additional storage savings from the new inline data compaction feature. When evaluating the storage efficiency of both inline compression and inline data compaction, we found savings of approximately 1.6:1 from this dataset.

Similarly, we conducted the same storage efficiency savings of data of files from a 1245GB customer production dataset. From this dataset, we measured a storage efficiency ratio of 1.7:1 from inline compression. The total storage saved from both inline compression and inline data compaction was 528GB.

As shown in Figure 2, both inline compression and inline data compaction represent additional storage efficiency savings in our test database configurations. The solid green section of each bar represents the amount of space saved by inline compression. The hashed green section represents the amount of space saved by inline data compaction and the blue section is the amount of data actually written to disk.

Figure 2) ONTAP 9 storage efficiency.




Space-efficient NetApp Snapshot copies can provide additional storage efficiency benefits. There is no performance penalty for creating Snapshot copies because data is never moved as it is with other copy-out technologies. The cost for Snapshot copies is only at the rate of block-level changes, not at 100% for each backup, as is the case with mirror copies. Snapshot technology can help you save on storage costs for backups and restores, and it opens up a number of efficient data management possibilities.

In a typical real-world SQL Server customer deployment, data volume Snapshot copies might be created every two hours and log and temp Snapshot copies every 15 minutes. These Snapshot copies can account for up to an additional 9:1 storage savings over a 24-hour period. Assuming 5% hourly overwrites and given our customer production database environment, the storage savings of each Snapshot copy would be 660GB, for a total of 7.6TB in savings. The total storage savings ratio achieved through inline compression, inline data compaction, and Snapshot copies was 15:1. Due to the unique header information that SQL Server adds to each block, deduplication provided little additional storage efficiency savings; therefore, we did not include calculations from this feature.

### 3.7 Performance Test Results

As shown in Figure 3, the SQL Server performance test results on ONTAP 9 demonstrated that the cluster IOPS increased from 4K IOPS at 266µs to a peak performance of 350K IOPS. For all load points up to 339K IOPS, we were able to maintain consistent storage latencies at or below 1ms. Similarly, we found that the SQL Server client latencies ranged from 370µs to 1470µs over this same interval. As shown in Figure 3, ONTAP 9 provides over 60% increase in IOPS under a millisecond over clustered Data ONTAP 8.3.1.

Figure 3) AFF8080A EX SQL Server 2014 performance.



## 4 Best Practices

NetApp recommends the following best practices:

- For the most IOPS-intensive OLTP databases, divide the database into multiple files striped across multiple volumes and LUNs (for block protocols).
- Set the FC HBA ports with a queue depth of 256.
- Set the load-balancing setting for the Microsoft DSM to least queue depth.

This list is not exhaustive. For more information, refer to Microsoft and NetApp documentation.

## 5 Conclusion

The NetApp AFF8080A EX running ONTAP 9.0 provides extremely high IOPS at consistently low latencies while serving a Microsoft SQL Server 2014 OLTP workload. Our testing showed that the AFF8080A EX cluster performed from 4K IOPs at 266µs to a peak performance of 350K IOPS at 1290µs. For all load points at or below 339K IOPS, we were able to maintain consistent storage latencies no greater than 1ms and provide total space-efficiency savings of up to 15:1 on customer production data.

Many of the financial and business advantages of AFF8080A EX derive from the extensive storage-efficiency portfolio available with NetApp storage solutions. NetApp storage-efficiency technologies work together on a single unified architecture, and they can be enabled or disabled to serve any requirement, application, or environment. NetApp leads the way in bringing value to its customers. The company has built its reputation based on storage efficiency, helping customers achieve what they previously thought impossible and partnering with customers to get the most value out of their IT environments.

## References

The following references are used in this TR:

- TR-4369: Best Practices Guide for Microsoft SQL Server and SnapManager 7.2 for SQL Server with Clustered Data ONTAP  
[www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-132372-16&m=TR-4369.pdf](http://www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-132372-16&m=TR-4369.pdf)
- SnapManager 7.2 for Microsoft SQL Server Installation and Setup Guide  
[https://library.netapp.com/ecm/ecm\\_get\\_file/ECMP11658051](https://library.netapp.com/ecm/ecm_get_file/ECMP11658051)
- Benchcraft TPC-E Toolkit  
<http://www.tpc.org/>

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 © 1994–2016 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, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Fitness, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, SnapCopy, Snap Creator, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-4516-0616