



Storage efficiency

ONTAP 9

NetApp
September 19, 2024

Table of Contents

- Storage efficiency 1
 - ONTAP storage efficiency overview 1
 - Thin provisioning 3
 - Deduplication 3
 - Compression 4
 - FlexClone volumes, files, and LUNs 5
 - Capacity measurements in System Manager 5
 - Temperature-sensitive storage efficiency overview 9
 - CPU or dedicated offload processor storage efficiency 10

Storage efficiency

ONTAP storage efficiency overview

Storage efficiency is the measure of how effectively a storage system utilizes available space by optimizing storage resources, minimizing wasted space, and reducing written data's physical footprint. Higher storage efficiency enables you to store the maximum amount of data within the smallest possible space at the lowest possible cost. For example, utilizing storage efficient technologies that detect and eliminate duplicate data blocks and data blocks filled with zeros decreases the overall amount of physical storage you need and reduces your overall cost.

ONTAP offers a wide range of storage efficiency technologies that reduce the amount of physical hardware or cloud storage consumed by your data and that also yield significant improvements to system performance, including faster reads of data, faster copies of datasets, and faster VM provisioning.

ONTAP storage efficiency technologies include:

- **Thin provisioning**

Thin provisioning enables you to allocate storage in a volume or LUN as it is needed instead of reserving it in advance. This reduces the amount of physical storage you need by allowing you to over-allocate your volumes or LUNs based on potential usage without reserving space that isn't currently being used.

- **Deduplication**

Deduplication reduces the amount of physical storage required for a volume in three distinct ways.

- **Zero block deduplication**

Zero block deduplication detects and eliminates data blocks filled with all zeros and only updates metadata. 100% of the space typically used by zero blocks is then saved. Zero block deduplication is enabled by default on all deduplicated volumes.

- **Inline deduplication**

Inline deduplication detects duplicate data blocks and replaces them with references to a unique shared block before data is written to disk. Inline deduplication speeds up VM provisioning by 20% to 30%. Depending upon your version of ONTAP and your platform, inline deduplication is available at the volume or aggregate level. It is enabled by default on AFF and ASA systems. You need to manually enable inline deduplication on FAS systems.

- **Background deduplication**

Background deduplication also detects duplicate data blocks and replaces them with references to a unique shared block, but further enhances storage efficiency by doing so after data is written to the disk. You can set up background deduplication to run when certain criteria are met on your storage system. For example, you might enable background deduplication to occur when your volume reaches 10% utilization. You can also manually trigger background deduplication or set it to run on a specific schedule. It is enabled by default on AFF and ASA systems. You need to manually enable background deduplication on FAS systems.

Deduplication is supported within volumes and across volumes within an aggregate. Reads of deduplicated

data typically incur no performance charge.

- **Compression**

[Compression](#) reduces the amount of physical storage required for a volume by combining data blocks in compression groups, each of which is stored as a single block. When a read or an overwrite request is received, only a small group of blocks is read, not the entire file. This process optimizes read and overwrite performance and enables greater scalability in the size of the files being compressed.

Compression can be run inline or postprocess. Inline compression provides immediate space savings by compressing data in memory before it is written to disk. Postprocess compression first writes the blocks to disk as uncompressed and then at a scheduled time compresses the data. It is enabled by default on AFA systems. You need to manually enable compression on all other systems.

- **Compaction**

Compaction reduces the amount of physical storage required for a volume by taking data chunks that are stored in 4 KB blocks, but that are less than 4 KB in size, and combining them into a single block. Compaction takes place while data is still in memory so unnecessary space is never consumed on disks. It is enabled by default on AFF and ASA systems. You need to manually enable compaction on FAS systems.

- **FlexClone volumes, files and LUNs**

[FlexClone technology](#) leverages Snapshot metadata to create writable, point-in-time copies of a volume, file or LUN. Copies share data blocks with their parents, consuming no storage except what is required for metadata until changes are written to a copy or its parent. When a change is written, only the delta is stored.

Where traditional dataset copies can take minutes or even hours to create, FlexClone technology lets you copy even the largest datasets almost instantaneously.

- **Temperature-sensitive storage efficiency**

ONTAP provides [temperature-sensitive storage efficiency](#) benefits by assessing how often your volume's data is accessed and mapping that frequency to the degree of compression applied to that data. For cold data that is accessed infrequently, larger data blocks are compressed. For hot data that is accessed frequently and is overwritten more often, smaller data blocks are compressed, making the process more efficient.

Temperature-sensitive storage efficiency (TSSE), introduced in ONTAP 9.8, is enabled automatically on newly created thinly provisioned AFF volumes. It is not enabled on [AFF A70, AFF A90, and AFF A1K platforms](#) that are introduced in ONTAP 9.15.1, which use a hardware offload processor.

- **CPU or dedicated offload processor storage efficiency**

Beginning with ONTAP 9.15.1, ONTAP provides [CPU or dedicated offload processor storage efficiency](#) and data compaction on AFF A70, AFF A90, and AFF A1K platforms. Storage efficiency is enabled automatically and requires no configuration.

You can realize the benefit of these technologies in your day-to-day operations with minimal effort. For example, suppose you need to supply 5,000 users with storage for home directories, and you estimate that the maximum space needed by any user is 1 GB. You could reserve a 5 TB aggregate in advance to meet the total potential storage need. However, you also know that home directory capacity requirements vary greatly across your organization. Instead of reserving 5 TB of total space for your organization, you can create a 2 TB

aggregate. Then you can use thin provisioning to nominally assign 1 GB of storage to each user but allocate the storage only as needed. You can actively monitor the aggregate over time and increase the actual physical size as necessary.

In another example, suppose you are using a Virtual Desktop Infrastructure (VDI) with a large amount of duplicate data among your virtual desktops. Deduplication reduces your storage usage by automatically eliminating duplicate blocks of information across the VDI, replacing them with a pointer to the original block. Other ONTAP storage efficiency technologies, such as compression, can also run in the background without your intervention.

ONTAP disk partitioning technology delivers greater storage efficiency as well. RAID DP technology protects against double disk failure without sacrificing performance or adding disk-mirroring overhead. Advanced SSD partitioning with ONTAP 9 increases usable capacity by almost 20%.

NetApp provides the same storage efficiency features available with on-premises ONTAP in the cloud. When you migrate data from on-premises ONTAP to the cloud, the existing storage efficiency is preserved. For example, suppose you have an SQL database containing business-critical data that you want to move from an on-premises system to the cloud. You can use data replication in BlueXP to migrate your data and, as part of the migration process, you can enable your latest on-premises policy for Snapshot copies in the cloud.

Thin provisioning

ONTAP offers a wide range of storage efficiency technologies in addition to Snapshot copies. Key technologies include thin provisioning, deduplication, compression, and FlexClone volumes, files, and LUNs. Like Snapshot copies, all are built on ONTAP's Write Anywhere File Layout (WAFL).

A *thin-provisioned* volume or LUN is one for which storage is not reserved in advance. Instead, storage is allocated dynamically, as it is needed. Free space is released back to the storage system when data in the volume or LUN is deleted.

Suppose that your organization needs to supply 5,000 users with storage for home directories. You estimate that the largest home directories will consume 1 GB of space.

In this situation, you could purchase 5 TB of physical storage. For each volume that stores a home directory, you would reserve enough space to satisfy the needs of the largest consumers.

As a practical matter, however, you also know that home directory capacity requirements vary greatly across your community. For every large user of storage, there are ten who consume little or no space.

Thin provisioning allows you to satisfy the needs of the large storage consumers without having to purchase storage you might never use. Since storage space is not allocated until it is consumed, you can “overcommit” an aggregate of 2 TB by nominally assigning a size of 1 GB to each of the 5,000 volumes the aggregate contains.


As long as you are correct that there is a 10:1 ratio of light to heavy users, and as long as you take an active role in monitoring free space on the aggregate, you can be confident that volume writes won't fail due to lack of space.

Deduplication

Deduplication reduces the amount of physical storage required for a volume (or all the

volumes in an AFF aggregate) by discarding duplicate blocks and replacing them with references to a single shared block. Reads of deduplicated data typically incur no performance charge. Writes incur a negligible charge except on overloaded nodes.

As data is written during normal use, WAFL uses a batch process to create a catalog of *block signatures*. After deduplication starts, ONTAP compares the signatures in the catalog to identify duplicate blocks. If a match exists, a byte-by-byte comparison is done to verify that the candidate blocks have not changed since the catalog was created. Only if all the bytes match is the duplicate block discarded and its disk space reclaimed.



Deduplication reduces the amount of physical storage required for a volume by discarding duplicate data blocks.

Compression

Compression reduces the amount of physical storage required for a volume by combining data blocks in *compression groups*, each of which is stored as a single block. Reads of compressed data are faster than in traditional compression methods because ONTAP decompresses only the compression groups that contain the requested data, not an entire file or LUN.

You can perform inline or postprocess compression, separately or in combination:

- *Inline compression* compresses data in memory before it is written to disk, significantly reducing the amount of write I/O to a volume, but potentially degrading write performance. Performance-intensive operations are deferred until the next postprocess compression operation, if any.
- *Postprocess compression* compresses data after it is written to disk, on the same schedule as deduplication.

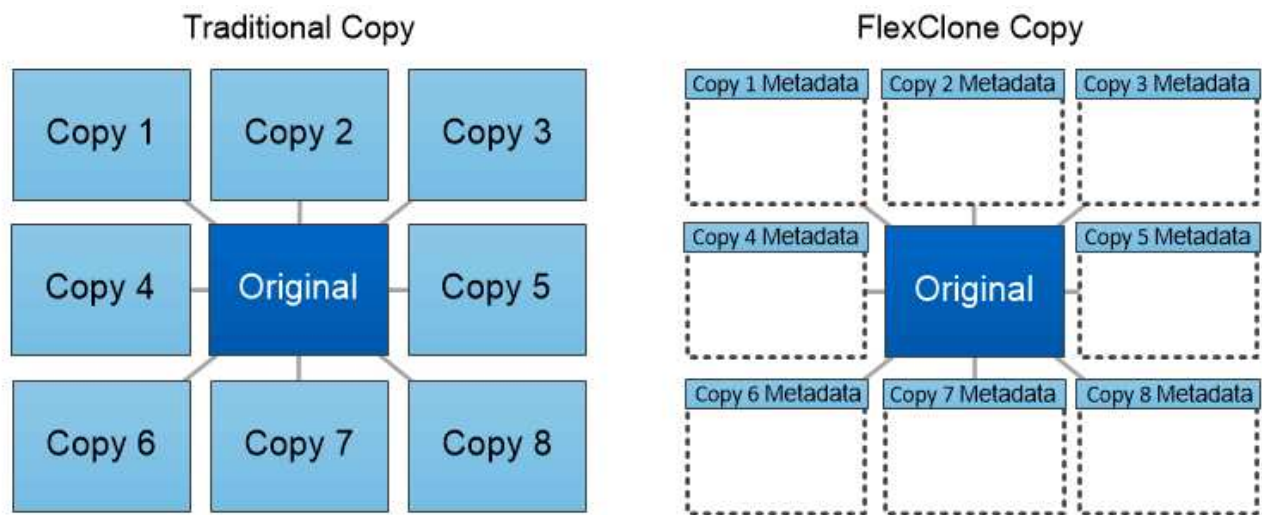
Inline data compaction Small files or I/O padded with zeros are stored in a 4 KB block whether or not they require 4 KB of physical storage. *Inline data compaction* combines data chunks that would ordinarily consume multiple 4 KB blocks into a single 4 KB block on disk. Compaction takes place while data is still in memory, so it is best suited to faster controllers.

FlexClone volumes, files, and LUNs

FlexClone technology references Snapshot metadata to create writable, point-in-time copies of a volume. Copies share data blocks with their parents, consuming no storage except what is required for metadata until changes are written to the copy. FlexClone files and FlexClone LUNs use identical technology, except that a backing Snapshot copy is not required.

Where traditional copies can take minutes or even hours to create, FlexClone software lets you copy even the largest datasets almost instantaneously. That makes it ideal for situations in which you need multiple copies of identical datasets (a virtual desktop deployment, for example) or temporary copies of a dataset (testing an application against a production dataset).

You can clone an existing FlexClone volume, clone a volume containing LUN clones, or clone mirror and vault data. You can *split* a FlexClone volume from its parent, in which case the copy is allocated its own storage.



FlexClone copies share data blocks with their parents, consuming no storage except what is required for metadata.

Capacity measurements in System Manager

System capacity can be measured as either physical space or logical space. Beginning with ONTAP 9.7, System Manager provides measurements of both physical and logical capacity.

The differences between the two measurements are explained in the following descriptions:

- **Physical capacity:** Physical space refers to the physical blocks of storage used in the volume or local tier. The value for physical used capacity is typically smaller than the value for logical used capacity due to the reduction of data from storage efficiency features (such as deduplication and compression).
- **Logical capacity:** Logical space refers to the usable space (the logical blocks) in a volume or local tier. Logical space refers to how theoretical space can be used, without accounting for results of deduplication or compression. The value for logical space used is derived from the amount of physical space used plus the savings from storage efficiency features (such as deduplication and compression) that have been

configured. This measurement often appears larger than the physical used capacity because it includes Snapshot copies, clones, and other components, and it does not reflect the data compression and other reductions in the physical space. Thus, the total logical capacity could be higher than the provisioned space.



In System Manager, capacity representations do not account for root storage tier (aggregate) capacities.

Measurements of used capacity

Measurements of used capacity are displayed differently depending on the version of System Manager you are using, as explained in the following table:

Version of System Manager	Term used for capacity	Type of capacity referred to
9.9.1 and later	Logical Used	Logical space used (if storage efficiency settings have been enabled)
9.7 and 9.8	Used	Logical space used (if storage efficiency settings have been enabled)
9.5 and 9.6 (Classic view)	Used	Physical space used

Capacity measurement terms

The following terms are used when describing capacity:

- **Allocated capacity:** The amount of space that has been allocated for volumes in a storage VM.
- **Available:** The amount of physical space available to store data or to provision volumes in a storage VM or on a local tier.
- **Capacity across volumes:** The sum of the used storage and available storage of all the volumes on a storage VM.
- **Client data:** The amount of space used by client data (either physical or logical).
 - Beginning with ONTAP 9.13.1, the capacity used by client data is referred to as **Logical Used**, and the capacity used by Snapshot copies is displayed separately.
 - In ONTAP 9.12.1 and earlier, the capacity used by client data added to the capacity used by Snapshot copies is referred to as **Logical Used**.
- **Committed:** The amount of committed capacity for a local tier.
- **Data reduction:**
 - Beginning with ONTAP 9.13.1, data reduction ratios are displayed as follows:
 - The data reduction value displayed on the **Capacity** panel is the ratio of logical used space compared to physical used space without considering the significant reductions achieved when using storage efficiency features, such as Snapshot copies.
 - When you display the details panel, you see both the ratio that was displayed on the overview

panel and the overall ratio of all logical used space compared to physical used space. Referred to as **With Snapshot copies**, this value includes the benefits derived from using Snapshot copies and other storage efficiency features.

- In ONTAP 9.12.1 and earlier, data reduction ratios are displayed as follows:
 - The data reduction value displayed on the **Capacity** panel is the overall ratio of all logical used space compared to physical used space, and it includes the benefits derived from using Snapshot copies and other storage efficiency features.
 - When you display the details panel, you see both the **Overall** ratio that was displayed on the overview panel and ratio of logical used space used only by client data compared to physical used space used only by client data, referred to as **Without Snapshot copies and clones**.

- **Logical used:**

- Beginning with ONTAP 9.13.1, the capacity used by client data is referred to as **Logical Used**, and the capacity used by Snapshot copies is displayed separately.
- In ONTAP 9.12.1 and earlier, the capacity used by client data added to capacity used by Snapshot copies is referred to as **Logical Used**.

- **Logical used %:** The percentage of the current logical used capacity compared to the provisioned size, excluding Snapshot reserves. This value can be greater than 100%, because it includes efficiency savings in the volume.
- **Maximum capacity:** The maximum amount of space allocated for volumes on a storage VM.
- **Physical used:** The amount of capacity used in the physical blocks of a volume or local tier.
- **Physical used %:** The percentage of capacity used in the physical blocks of a volume compared to the provisioned size.
- **Provisioned capacity:** A file system (volume) that has been allocated from a Cloud Volumes ONTAP system and is ready to store user or application data.
- **Reserved:** The amount of space reserved for already provisioned volumes in a local tier.
- **Used:** The amount of space that contains data.
- **Used and reserved:** The sum of physical used and reserved space.

Capacity of a storage VM

The maximum capacity of a storage VM is determined by the total allocated space for volumes plus the remaining unallocated space.

- The allocated space for volumes is the sum of the used capacity and the sum of available capacity of FlexVol volumes, FlexGroup volumes, and FlexCache volumes.
- The capacity of volumes is included in the sums, even when they are restricted, offline, or in the recovery queue after deletion.
- If volumes are configured with auto-grow, the maximum autosize value of the volume is used in the sums. Without auto-grow, the actual capacity of the volume is used in the sums.

The following chart explains how the measurement of the capacity across volumes relates to the maximum capacity limit.



Beginning with ONTAP 9.13.1, cluster administrators can [enable a maximum capacity limit for a storage VM](#). However, storage limits cannot be set for a storage VM that contains volumes that are for data protection, in a SnapMirror relationship, or in a MetroCluster configuration. Also, quotas cannot be configured to exceed the maximum capacity of a storage VM.

After the maximum capacity limit is set, it cannot be changed to a size that is less than the currently allocated capacity.

When a storage VM reaches its maximum capacity limit, certain operations cannot be performed. System Manager provides suggestions for next steps in [Insights](#).

Capacity measurement units

System Manager calculates storage capacity based on binary units of 1024 (2^{10}) bytes.

- Beginning with ONTAP 9.10.1, storage capacity units are displayed in System Manager as KiB, MiB, GiB, TiB, and PiB.
- In ONTAP 9.10.0 and earlier, these units are displayed in System Manager as KB, MB, GB, TB, and PB.



The units used in System Manager for throughput continue to be KB/s, MB/s, GB/s, TB/s, and PB/s for all releases of ONTAP.

Capacity unit displayed in System Manager for ONTAP 9.10.0 and earlier	Capacity unit displayed in System Manager for ONTAP 9.10.1 and later	Calculation	Value in bytes
KB	KiB	1024	1024 bytes
MB	MiB	1024 * 1024	1,048,576 bytes
GB	GiB	1024 * 1024 * 1024	1,073,741,824 bytes
TB	TiB	1024 * 1024 * 1024 * 1024	1,099,511,627,776 bytes
PB	PiB	1024 * 1024 * 1024 * 1024 * 1024	1,125,899,906,842,624 bytes

Related information

[Monitor capacity in System Manager](#)

[Logical space reporting and enforcement for volumes](#)

Temperature-sensitive storage efficiency overview

ONTAP provides temperature-sensitive storage efficiency benefits by assessing how often your volume's data is accessed and mapping that frequency to the degree of compression applied to that data. For cold data that is accessed infrequently, larger data blocks are compressed, and for hot data, which is accessed frequently and is overwritten more often, smaller data blocks are compressed, making the process more efficient.

Temperature-sensitive storage efficiency (TSSE) is introduced in ONTAP 9.8 and is enabled automatically on newly created thinly provisioned AFF volumes. You can enable temperature-sensitive storage efficiency on existing AFF volumes and on thinly provisioned non-AFF DP volumes.

Introduction of "default" and "efficient" modes

Beginning with ONTAP 9.10.1, *default* and *efficient* volume-level storage efficiency modes are introduced for AFF systems only. The two modes provide a choice between file compression (default), which is the default mode when creating new AFF volumes, or temperature-sensitive storage efficiency (efficient), which enables temperature-sensitive storage efficiency. With ONTAP 9.10.1, [temperature-sensitive storage efficiency must be explicitly set](#) to enable auto adaptive compression. However, other storage efficiency features like data-compaction, auto dedupe schedule, inline deduplication, cross volume inline deduplication, and cross volume background deduplication are enabled by default on AFF platforms for both default and efficient modes.

Both storage efficiency modes (default and efficient) are supported on FabricPool-enabled aggregates and with all tiering policy types.

Temperature-sensitive storage efficiency enabled on C-Series platforms

Temperature-sensitive storage efficiency is enabled by default on AFF C-Series platforms and when migrating volumes from a non-TSSE platform to a TSSE-enabled C-Series platform using volume move or SnapMirror with the following releases installed on the destination:

- ONTAP 9.12.1P4 and later
- ONTAP 9.13.1 and later

For more information, see [Storage efficiency behavior with volume move and SnapMirror operations](#).

For existing volumes, temperature-sensitive storage efficiency is not enabled automatically, however, you can [modify the storage efficiency mode](#) manually to change to efficient mode.



Once you change the storage efficiency mode to efficient you cannot change it back.

Improved storage efficiency with sequential packing of contiguous physical blocks

Beginning with ONTAP 9.13.1, temperature-sensitive storage efficiency adds sequential packing of contiguous physical blocks to further improve storage efficiency. Volumes that have temperature-sensitive storage efficiency enabled automatically have sequential packing enabled when you upgrade systems to ONTAP 9.13.1. After sequential packing is enabled, you must [manually repack existing data](#).

Upgrade considerations

When upgrading to ONTAP 9.10.1 and later, existing volumes are assigned a storage efficiency mode based on the type of compression currently enabled on the volumes. During an upgrade, volumes with compression enabled are assigned the default mode, and volumes with temperature-sensitive storage efficiency enabled are assigned the efficient mode. If compression is not enabled, storage efficiency mode remains blank.

CPU or dedicated offload processor storage efficiency

Beginning with ONTAP 9.15.1, ONTAP provides storage efficiency and data compaction on AFF A70, AFF A90, and AFF A1K platforms. Depending on the platform, compression is performed using either the main CPU or with a dedicated offload processor. Storage efficiency is enabled automatically and requires no configuration. Storage efficiency is enabled by default on newly created thin-provisioned volumes and is applied to existing data, including volumes moved from other platforms to AFF A70, AFF A90, or AFF A1K platforms.

Temperature-sensitive storage efficiency is not applied on AFF A70, AFF A90, and AFF A1K platforms. Compression is not based on hot or cold data on these platforms, so compression begins without waiting for data to become cold.

Storage efficiency on AFF A70, AFF A90, and AFF A1K platforms uses sequential packing of contiguous physical blocks to further improve storage efficiency for compressed data.

Data that's migrated to AFF A70, AFF A90, or AFF A1K platforms using volume move or SnapMirror technology is converted automatically to 32k inline compression. Storage efficiency for data migrated from an AFF A70, A90, or A1K platform to an earlier platform is transformed automatically based on the destination platform's efficiency capabilities.

For information about upgrading a controller to an AFF A70, AFF A90, or AFF A1K, see the [ONTAP Hardware Upgrade Documentation](#).

Copyright information

Copyright © 2024 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, nonsublicensable, 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.