



Create storage resources

Workload Factory for NetApp Console

NetApp

October 07, 2026

Table of Contents

- Create storage resources. 1
 - Create or discover cloud storage systems in NetApp Console 1
 - Create a storage system 1
 - Create an FSx for ONTAP file system 1
 - Discover an existing cloud storage system 2
 - Create a storage VM for an FSx for ONTAP file system 3
 - Create a storage VM 3
 - Create an FSx for ONTAP volume in Workload Factory 3
 - About this task 4
 - Before you begin 5
 - Create a volume. 5
 - Create an initiator group for a file system in NetApp Workload Factory 9
 - Create a block device for a file system in NetApp Workload Factory 10

Create storage resources

Create or discover cloud storage systems in NetApp Console

Create or discover cloud storage systems to add and manage volumes and additional data services from the NetApp Console.

Create a storage system

The first step is to create a storage system. If you already created a storage system in your cloud provider's console, you can [discover it using the NetApp Console](#).

About this task

A storage VM is created when you create a file system.

Before you begin

Before creating your FSx for ONTAP file system, you will need:

- The ARN of an IAM role that gives Workload Factory the permissions needed to create an FSx for ONTAP file system. [Learn how to grant permissions to an AWS account](#).
- The region and VPC information for where you will create the FSx for ONTAP instance.

Create an FSx for ONTAP file system

You can create an FSx for ONTAP file system using *Quick create* or *Advanced create*. You can also use the following tools available in the Codebox: REST API, CloudFormation, and Terraform. [Learn how to use Codebox for automation](#).



When using Terraform from Codebox, the code you copy or download hides `fsxadmin` and `vsadmin` passwords. You'll need to re-enter the passwords when you run the code.

Quick create

Quick create enables you to use a recommended best-practice configuration. You can change most settings after you create an FSx for ONTAP file system.

Steps

1. From the NetApp Console menu, select **Storage** and then **Management**.
2. Select **Add system** from the Systems page.
3. Select **Amazon Web Services** as the location, and then select **Add new** for Amazon FSx for NetApp ONTAP.
Unresolved directive in fsx-storage/create-discover-systems.adoc - include::_include/create-file-system-quick.adoc[]

Result

If you created the file system, the new FSx for ONTAP configuration appears on the Systems page.

You can manage your FSx for ONTAP file systems in several ways, such as from Workloads in the

NetApp Console, using ONTAP System Manager, and using AWS CloudFormation.

Advanced create

With Advanced create, you set all of the configuration options, including availability, security, backups, and maintenance.

Steps

1. From the NetApp Console menu, select **Storage** and then **Management**.
2. Select **Add system** from the Systems page.
3. Select **Amazon Web Services** as the location, and then select **Add new** for Amazon FSx for NetApp ONTAP.

Unresolved directive in fsx-storage/create-discover-systems.adoc - include::_include/create-file-system-advanced.adoc[]

Result

If you created the file system, the new FSx for ONTAP configuration appears on the Systems page.

You can manage your FSx for ONTAP file systems in several ways, such as from Workloads in the NetApp Console, using ONTAP System Manager, and using AWS CloudFormation.

Discover an existing cloud storage system

If you previously provided your credentials for your cloud storage provider in Workload Factory, you can automatically discover your storage systems from the Discoverable systems page. You can also review available data services.

About this task

You can discover a cloud storage system only once within an account and attach it to one workspace. The storage system can later be removed and re-associated to a different workspace.

Steps

1. From the NetApp Console menu, select **Storage**, then **Management**, and then **Discoverable systems**.
2. The count of discovered cloud storage systems displays. Select **Discover**.
3. Select one or more storage systems and select **Discover** to add them to the Systems page.



- If you select an un-named storage system, you will receive a prompt to enter a name for the system.
- If you select a storage system that doesn't have the credentials required to manage it from the Console, you'll receive a prompt to select credentials with the required permissions.
- The following regions aren't supported for discovery: China regions, GovCloud (US) regions, Secret Cloud, and Top Secret Cloud.

Result

The Console displays your discovered system on the Systems page. You can manage your storage systems in several ways, such as from Workloads in the NetApp Console and using ONTAP System Manager or using your cloud provider's storage management tools.

Create a storage VM for an FSx for ONTAP file system

Create a storage VM (SVM) for an FSx for ONTAP file system to access storage and data services virtually for your workloads in NetApp Workload Factory.

About this task

Storage VMs are isolated file servers that you can use to access the data from each workload in Workload Factory Storage. Each SVM has its own administrative credentials and endpoints for administering and accessing data.

With SVMs, when you access data in FSx for ONTAP, your clients and workstations mount a volume, CIFS/SMB share, or iSCSI LUN hosted by an SVM using the SVM's endpoint (IP address).

Before you begin

Verify the supported number of storage VMs per file system. Refer to [Managing FSx for ONTAP storage virtual machines](#) in AWS documentation for the maximum number of SVMs per file system.

Create a storage VM

You can create a storage VM from the Workload Factory console. You can also use the following tools available in the Codebox: REST API, CloudFormation, and Terraform. [Learn how to use Codebox for automation.](#)



When using Terraform from Codebox, the code you copy or download hides `fsxadmin` and `vsadmin` passwords. You'll need to re-enter the passwords when you run the code.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Storage** > **FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system and then select **Manage**.
4. In the file system overview under Storage VMs, select **Create storage VM**.
5. On the Create storage VM page, under Storage VM configuration, provide the following:
 - a. **Name**: Enter a name for the storage VM.
 - b. **Storage VM credentials**: Provide a password for this storage VM's `vsadmin` user or use the file system's `fsxadmin` user credentials.
 - c. **Root volume security style**: Select the root volume security style depending on the type of clients that access your data - UNIX (Linux clients), NTFS (Windows clients), or Mixed.
 - d. **Tags**: Optionally, you can add up to 50 tags.
6. Select **Create**.

Create an FSx for ONTAP volume in Workload Factory

After setting up your FSx for ONTAP file system, create FSx for ONTAP volumes in Workload Factory as virtual resources for grouping your data.

About this task

FSx for ONTAP volumes group data virtually, determine how data is stored, and determine the type of access to your data. Volumes don't consume file system storage capacity. The data that is stored in a volume primarily consumes SSD storage. Depending on the volume's tiering policy, the data might also consume capacity pool storage. You set a volume's size when you create it, and you can change its size later.

The following protocols might be used for your volumes:

- SMB/CIFS: file storage protocol for Windows operating systems
- NFS: file storage protocol for Unix operating systems
- iSCSI: block storage protocol

S3 endpoints can be attached to an FSx for ONTAP volume. Using an S3 access point, you can access file data residing on SMB/CIFS or NFS volumes via the AWS S3 APIs. This allows you to integrate your existing data with GenAI, ML, and analytics from AWS services that support S3 access points.

Details for volume settings

Immutable files

This feature, also known as SnapLock, is disabled by default. Enabling immutable files prevents data deletion or overwriting for a set period. Enabling this feature is possible only during volume creation. After the feature is enabled, it cannot be disabled. This is a premium feature for FSx for ONTAP that carries an additional charge. For more information, refer to [How SnapLock works](#) in Amazon FSx for NetApp ONTAP documentation.

- **Retention modes:** You can select from two retention modes - *Enterprise* or *Compliance*.
 - In *Enterprise* mode, an immutable files, or SnapLock, administrator can delete a file during its retention period.
 - In *Compliance* mode, a WORM file cannot be deleted before its retention period expires. Similarly, the immutable volume cannot be deleted until the retention periods for all files within the volume expire.
- **Retention period:** The retention period has two settings - *retention policy* and *retention periods*. The *retention policy* defines how long to retain files in an immutable WORM state. You can specify your own retention policy or use the default retention policy (unspecified), which is 30 years. The minimum and maximum *retention periods* define the range of time allowed for locking files.



Even after the retention period expires, you can't modify a WORM file. You can only delete it or set a new retention period to turn on WORM protection again.

- **Autocommit:** You'll have the option to enable the autocommit feature. The autocommit feature commits a file to WORM state on a SnapLock volume if the file did not change for the autocommit period duration. The autocommit feature is disabled by default. You must ensure that the files you want to autocommit reside on a SnapLock volume.
- **Privileged delete:** A SnapLock administrator can turn on privileged delete on a SnapLock Enterprise volume to allow a file to be deleted before the file's retention period expires. This feature is disabled by default.
- **Volume append mode:** You can't modify existing data in a WORM-protected file. However, immutable files allows you to maintain protection for existing data using WORM-appendable files. For example, you can generate log files or preserve audio or video streaming data while writing data to them incrementally. [Learn more about volume-append mode](#) in Amazon FSx for NetApp ONTAP documentation.

Before you begin

Review the following prerequisites before you create a volume:

- You must have an FSx for ONTAP file system in the Workload Factory console.
- You must have a storage VM.
- For protocol access, complete the following:
 - To configure access to the volume, you must associate a link. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
 - You must configure access for the protocol you select, either SMB/CIFS, NFS, or iSCSI.

Create a volume

You can create a volume using the following tools available in the Codebox: REST API, CloudFormation, and Terraform. [Learn how to use Codebox for automation](#).



When using Terraform from Codebox, the code you copy or download hides `fsxadmin` and `vsadmin` passwords. You'll need to re-enter the passwords when you run the code.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select **Create volume**.
4. On the Create volume page under General details, provide the following details:
 - a. **Volume name**: Enter a name for the volume.
 - b. **Storage VM name**: Optionally, enter a storage VM name.
 - c. **Volume style**: Select **FlexVol** or **FlexGroup** volume.
 - A *FlexVol volume* is a type of logical storage in ONTAP that allows you to manage data. FlexVol volumes can be moved between aggregates or nodes within the same storage virtual machine (SVM) without disrupting client access. The FlexVol volume style is selected by default.
 - A *FlexGroup* volume is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. It consists of several member volumes (constituents), which are underlying FlexVol volumes that share traffic transparently.
 - d. **Volume size**: Enter the volume size and unit.

Optionally, enable volume autogrow. This option is available when you select **File access** as the volume access type.

+

FlexGroup volume size depends on the number of constituents, requiring 100 GiB per constituent.

- e. **Volume autogrow**: Optionally, enable volume autogrow to automatically expand volume capacity until the volume reaches the maximum size. This feature accommodates increasing data usage, ensuring uninterrupted operations.

Specify the maximum volume growth size and unit. You cannot set the maximum growth size smaller than the current volume size

f. **Tags:** Optionally, you can add up to 50 tags.

5. Under Access (only for file systems with associated links), provide the following details:

a. **Access type:** Select **File access** or **Block access**. Additional fields to configure volume access differ depending on your selection.

- **File access:** allows multiple authorized users and devices access to the volume using SMB/CIFS, NFS, or dual (SMB/NFS) protocols.

Complete the following fields to set up file access to the volume.

b. **NFS export policy:** Provide the following details to provide NFS access:

i. **Access control:** Select a **Custom export policy**, **Existing export policy**, or **No access to the volume** from the dropdown menu.

ii. **Export policy name:**

If you selected a custom export policy, select an existing policy name from the dropdown menu.

If you selected an existing export policy, enter a new policy name.

iii. **Add Export Policy Rule:** Optionally, for a custom export policy, you can add export policy rules to the policy.

c. **SMB/CIFS share:** Provide the following:

i. **Name:** Enter the SMB/CIFS share name to provide access.

ii. **Permissions:** Select Full control, Read/Write, Read, or No access, and then enter the users or groups separated by a semicolon (;). Users or groups are case sensitive and the user's domain must be included using the format "domain\username".

d. **Security style:** For dual-protocol volumes, select either the UNIX or NTFS security style. UNIX is the default security style for dual-protocol volumes. For detailed guidance on user mapping in this context, refer to the AWS blog article ["Enabling multiprotocol workloads with Amazon FSx for NetApp ONTAP"](#).

- **Block access:** allows hosts running critical business applications access to the volume using the iSCSI protocol. Block access is only available when the file system scale-out deployment has six HA pairs or fewer.

Complete the following fields to set up block access to the volume.

i. **iSCSI configuration:** Provide the following details to configure iSCSI for block access to the volume.

A. Select **Create a new initiator group** or **Map an existing initiator group**.

B. Select the **Host operating system** from the dropdown menu.

C. Enter an **Initiator group name** for a new initiator group.

D. Under Host Initiators, add one or more iSCSI qualified name (IQN) host initiators.

e. **S3 access point:** Optionally, attach an S3 access point to access FSx for ONTAP file system data residing on NFS or SMB/CIFS volumes via AWS S3 APIs. Only the file access type is supported. Providing the following details:

- **S3 access point name:** Enter the name of the S3 access point.
- **User:** Select an existing user with access to the volume or create a new user.
- **User type:** Select **UNIX** or **Windows** as the user type.

- **Network configuration:** Select **Internet** or **Virtual private cloud (VPC)**. The type of network you choose determines whether the access point is accessible from the internet or restricted to a specific VPC.
- **Enable inventory table:** When you enable the inventory table on the volume, the system generates metadata for all objects accessible to the S3 access point and incurs AWS S3 request costs. Refer to [Amazon S3 pricing documentation](#) for more information.

f. **S3 access point tags:** Optionally, you can add up to 50 tags or remove tags.

6. Under Efficiency and protection, provide the following details:

a. **Storage efficiency:** Enabled by default. Select to disable the feature.

ONTAP achieves storage efficiency using deduplication and compression features. Deduplication eliminates duplicate data blocks. Data compression compresses the data blocks to reduce the amount of physical storage that is required.

b. **Snapshot policy:** Select the snapshot policy to specify the frequency and retention of snapshots.

The following are default policies from AWS. To display existing snapshot policies, you must [associate a link](#).

default

This policy automatically creates snapshots on the following schedule, with the oldest snapshot copies deleted to make room for newer copies:

- A maximum of six hourly snapshots taken five minutes past the hour.
- A maximum of two daily snapshots taken Monday through Saturday at 10 minutes after midnight.
- A maximum of two weekly snapshots taken every Sunday at 15 minutes after midnight.



Snapshot times are based on the file system's time zone, which defaults to Coordinated Universal Time (UTC). For information about changing the time zone, refer to [Displaying and setting the system time zone](#) in the NetApp Support documentation.

default-1weekly

This policy works in the same way as the `default` policy, except that it only retains one snapshot from the weekly schedule.

none

This policy doesn't take any snapshots. You can assign this policy to volumes to prevent automatic snapshots from being taken.

c. **Tiering policy:** Select the tiering policy for the data stored in the volume.

Balanced (Auto) is the default tiering policy when creating a volume using the Workload Factory console. For more information about volume tiering policies, refer to [Volume storage capacity](#) in AWS FSx for NetApp ONTAP documentation. Note that Workload Factory uses use-case based names in the Workload Factory console for tiering policies and includes FSx for ONTAP tiering policy names in parentheses.

d. **Immutable files:** Enabling the immutable files feature permanently commits files in this volume to an

immutable WORM (write-once-read-many) state. Provide the following details:

- i. Select to enable **Immutable files powered by SnapLock**.
- ii. Select the box to agree and proceed.
- iii. Select **Enable**.
- iv. **Retention mode**: Select **Enterprise** or **Compliance** mode.
- v. **Retention period**:
 - Select the retention policy:
 - **Unspecified**: Sets the retention policy to 30 years.
 - **Specify period**: Enter the number of seconds, minutes, hours, days, months, or years to set your own retention policy.
 - Select the minimum and maximum retention periods:
 - **Minimum**: Enter the number of seconds, minutes, hours, days, months, or years to set the minimum retention period.
 - **Maximum**: Enter the number of seconds, minutes, hours, days, months, or years to set the maximum retention period.
- vi. **Autocommit**: Disable or enable autocommit. If you enable autocommit, set the autocommit period.
- vii. **Privileged delete**: Disable or enable. If you enable privileged delete, a SnapLock administrator can delete a file before its retention period expires.
- viii. **Volume append mode**: Disable or enable. Enables you to add new content to WORM files.
- e. **ARP/AI**: NetApp Autonomous Ransomware Protection with AI (ARP/AI) is enabled by default when a link is associated with the file system. [Learn more about ARP/AI](#). Accept the statement to proceed.

If the feature is unavailable, it is because of one of the following reasons:

- A link is not associated with the file system. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
- Volumes with immutable files and volumes with the NVMe protocol are not supported for ARP/AI. iSCSI volumes are supported beginning with ONTAP 9.17.1.
- The file system already has an ARP/AI policy.

7. Under Advance configuration, enter the following details. The fields under this section differ based on the volume style you select.
 - a. **Junction path**: Enter the location in the storage VM's namespace where the volume gets mounted. The default junction path is `/<volume-name>`.
 - b. **Aggregate name**: Only for FlexVol volumes. Select the aggregate to host the new volume.
 - c. **Aggregates list**: Only for FlexGroup volumes. Add or remove aggregates. The minimum number of aggregates is one.
 - d. **Number of constituents**: Only for FlexGroup volumes. Enter the number of constituents per aggregate. 100 GiB is required per constituent.
8. Select **Create**.

Related information

- [Adjust volume capacity in Workload Factory](#)
- [Change volume tiering policy in Workload Factory](#)

- [Manage S3 access points in Workload Factory](#)

Create an initiator group for a file system in NetApp Workload Factory

Use NetApp Workload Factory to create initiator groups and manage host access to SAN block devices.

About this task

Initiator groups, or igroups, connect block devices (LUNs) to the compute resources that are allowed to access them. Unlike NFS or CIFS, where a volume is broadly accessible and user permissions control access, block storage permissions operate at the machine level. Typically, only one system can access a block device at a time.

An igroup acts as a permission layer for block storage. When a server connects to the storage system, it identifies itself using its iSCSI qualified (IQN) host initiator. If that IQN belongs to one or more igroups, then the server gains access to all LUNs associated with those igroups. Both an igroup and an iSCSI host connection are required for iSCSI to function properly.

Before you begin

You must associate a link to create igroups. [Learn how to associate an existing link or to create and associate a new link](#). After you associate the link, return to this operation.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the FSx for ONTAP page, select the **Block devices** tab.
4. Select the resource type **Create initiator group** and then select **Create igroup**.
5. In the **Create initiator group** dialog, do the following:
 - **igroup name**: Enter a name for the initiator group.
 - **igroup description**: (Optional) Enter a description for the initiator group.
 - **Storage VM name**: Select the storage VM for the initiator group.
 - **Block device name**: Select one or more block devices to associate with the initiator group. The block devices listed are those that have not been mapped to a host initiator yet.
 - **Operating system type**: Select Linux, VMware, or Windows for the operating system type.
 - **Host initiators**: Add one or more iSCSI qualified (IQN) host initiators to the initiator group.
6. Select **Create**.

Related information

[Manage the block storage for an FSx for ONTAP file system](#)

Create a block device for a file system in NetApp Workload Factory

Create block devices to support your line of business (LOB) application requirements. You can create block devices in an existing volume, create a new volume for the new block device, or create a new volume from a snapshot for the new block device.

About this task

Only FlexVol volumes are supported for block devices in NetApp Workload Factory. You can create block devices using the iSCSI protocol.

The block size must be smaller than the available FlexVol volume size.

Before you begin

- You must associate a link to create block devices. [Learn how to associate an existing link or to create and associate a new link](#). After you associate the link, return to this operation.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Storage** > **FSx for ONTAP**.
3. From FSx for ONTAP, select the file system to create a block device for, and then select the **Block devices** tab.
4. Select **Create block device**.
5. Under **Volume details**, do the following:
 - a. **Volume name**: Select one of the following options.
 - Create a new volume and enter the name of the volume.
 - Select an existing volume.
 - Create from a volume snapshot. Select the source volume name and snapshot name. Only volumes without block devices are available.
 - b. **Aggregate name**: Select the aggregate to host the new volume.
 - c. **Storage VM**: Select a storage VM.
 - d. **Volume style**: The default volume style is **FlexVol**.
 - e. **Volume size**: Enter the size of the volume and select the unit. The maximum size per FlexVol volume is 100 TiB.
 - f. **Volume autogrow**: Optionally, enable volume autogrow to allow the volume to automatically increase in size when it reaches capacity. The maximum growth size is 300 TiB.
 - g. **Tags**: Optionally, add tags to help organize and categorize your block device.
6. Under **Block device details**, do the following:
 - a. **Block device name**: Enter a name for the block device.
 - b. **Block device size**: Enter the size of the block device and select the unit. The block device size must be smaller than the available volume size.
7. Under **Access**, do the following:

a. **iSCSI configuration:** Select one of the following options.

- **Create new initiator group:** Provide the host operating system, initiator group name, and add one or more iSCSI qualified name (IQN) host initiators.
- **Map existing initiator group:** Select an existing initiator group, provide the host operating system, and select one or more iSCSI qualified name (IQN) host initiators.

8. Under **Efficiency and protection**, do the following:

a. **Storage efficiency:** Enabled by default. Select to disable the feature.

ONTAP achieves storage efficiency using deduplication and compression features. Deduplication eliminates duplicate data blocks. Data compression compresses the data blocks to reduce the amount of physical storage that is required.

b. **Snapshot policy:** Select the snapshot policy to specify the frequency and retention of snapshots.

The following are default policies from AWS. To display existing snapshot policies, you must [associate a link](#).

default

This policy automatically creates snapshots on the following schedule, with the oldest snapshot copies deleted to make room for newer copies:

- A maximum of six hourly snapshots taken five minutes past the hour.
- A maximum of two daily snapshots taken Monday through Saturday at 10 minutes after midnight.
- A maximum of two weekly snapshots taken every Sunday at 15 minutes after midnight.



Snapshot times are based on the file system's time zone, which defaults to Coordinated Universal Time (UTC). For information about changing the time zone, refer to [Displaying and setting the system time zone](#) in the NetApp Support documentation.

default-1weekly

This policy works in the same way as the `default` policy, except that it only retains one snapshot from the weekly schedule.

none

This policy doesn't take any snapshots. You can assign this policy to volumes to prevent automatic snapshots from being taken.

c. **Tiering policy:** Select the tiering policy for the data stored in the volume.

Balanced (Auto) is the default tiering policy when creating a volume using the Workload Factory console. For more information about volume tiering policies, refer to [Volume storage capacity](#) in AWS FSx for NetApp ONTAP documentation. Note that Workload Factory uses use-case based names in the Workload Factory console for tiering policies and includes FSx for ONTAP tiering policy names in parentheses.

d. **ARP/AI:** NetApp Autonomous Ransomware Protection with AI (ARP/AI) is enabled by default when a link is associated with the file system. [Learn more about ARP/AI](#). Accept the statement to proceed.

If the feature is unavailable, it is because of one of the following reasons:

- A link is not associated with the file system. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
- Volumes with immutable files and volumes with the NVMe protocol are not supported for ARP/Al.
- The file system already has an ARP/Al policy.

9. Select **Create**.

Related information

[Manage the block storage for an FSx for ONTAP file system](#)

Copyright information

Copyright © 2026 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.