



Manage S3 access points

Workload Factory for NetApp Console

NetApp
October 07, 2026

Table of Contents

- Manage S3 access points 1
 - Create S3 access points for a volume in NetApp Workload Factory 1
 - Manage the S3 access points for a volume in NetApp Workload Factory 2
 - Edit the S3 access points for a volume 2
 - View S3 access point details for a volume 2
 - Delete the S3 access points for a volume in NetApp Workload Factory 3
- Set up the Journal table infrastructure for NetApp Workload Factory 3
 - About this task 4
 - Before you begin 4
 - Set up the journal table infrastructure 7
 - Troubleshoot the infrastructure setup for the journal table 7
 - Permissions reference for the journal table setup 10

Manage S3 access points

Create S3 access points for a volume in NetApp Workload Factory

Create and attach S3 access points for a volume in NetApp Workload Factory.

About this task

Amazon FSx for NetApp ONTAP lets NFS and SMB file systems access S3 data and connect to AWS services like Amazon Bedrock, SageMaker, Athena, AWS Glue, and more. You can connect AWS services to all of your object storage data.

Attach S3 access points to NFS and SMB volumes in an FSx for ONTAP file system so AWS services access files as if they are in an S3 bucket. When attaching the access point, define a unique ID, select the file access type (UNIX or Windows), and add a username for access authorization.

After you attach the S3 access point, it appears in the AWS Management Console with a unique alias. Use this alias as the S3 bucket name for AWS services, such as Amazon Bedrock, to access files in the FSx for ONTAP volume.

You can attach multiple S3 access points to a single FSx for ONTAP volume, each with a unique access level, to connect to multiple AWS services.

Before you begin

Complete the following requirements before you begin.

- You must have an existing volume with an S3 access point. [Create a volume with an S3 access point](#)
- You must [grant credentials with the *operations and remediation* permission policy](#) in Workload Factory to complete this task.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume you want to use, then select **Advanced actions**, and then **Manage S3 access points**.
4. On the Manage S3 access points page, select **Create and attach S3 access point**.
5. On the **Create and attach S3 access point** page, provide the following information:
 - **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. To use the Journal table feature, you must select **Internet** as the network configuration.
 - **Inventory table**: Optional. Requires the `getObject` permission. Enable the inventory table on the volume to generate 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. The table updates

every 24 hours.

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

6. Select **Create**.

Related information

Optionally, you can use the Journal table feature with S3 access points in Workload Factory. The Journal table infrastructure captures and stores audit logs of user access events and object operations across Amazon FSx for ONTAP volume access points. To enable the Journal table feature, you must set up the necessary AWS infrastructure, including AWS CloudTrail, AWS CloudWatch, AWS S3 Buckets, AWS CloudWatch log groups, and AWS Identity and Access Management (IAM) roles and policies. [Set up the journal table infrastructure for NetApp Workload Factory](#).

Other related topics:

[Edit or delete S3 access points for a volume in NetApp Workload Factory](#)

Manage the S3 access points for a volume in NetApp Workload Factory

Update S3 access point details or delete S3 access points for a volume in NetApp Workload Factory.

You can also enable or disable the **Inventory table** for the access point. The inventory table 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.

Edit the S3 access points for a volume

Edit the S3 access points for a volume in NetApp Workload Factory.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > Volumes explorer**.
3. Select the actions menu for the volume you want to edit S3 access points for, then select **Advanced actions**, and then **Manage S3 access points**.
4. From the **Manage S3 access points** screen, select the actions menu and then select **Edit access point**.
5. Make updates and then select **Apply**.

To set up the infrastructure for the journal table feature, you must select **Internet** as the network configuration for the S3 access point. [Learn how to set up the Journal table infrastructure for NetApp Workload Factory](#).

View S3 access point details for a volume

View S3 access point details like the alias, ARN, and S3 URI. With inventory and journal tables enabled, you can view metadata and audit logs of user access events and object operations across S3 access points.

About this task

When the inventory table is enabled on the volume, you can view access point, inventory table, and table bucket details of existing S3 access points attached to the volume. The system also provides a link to the inventory table in the AWS Management Console.

When the journal table is enabled, Workload Factory generates logs on access points in a table for your reference. The table includes CloudTrail events, trails (a record of AWS activities), CloudTrail filters, CloudWatch log groups, and more.

You can copy access point details for use in other applications.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > Volumes explorer**.
3. Select the actions menu for the volume to view details for, then select **Advanced actions**, and then **View details**.

Details appear under four tabs: **Access point**, **Inventory table**, **Journal table**, and **Table bucket**.
4. If the status for inventory or journal tables is enabled, but you're missing the `get:Object` permission, follow the instructions in the dialog to add the required permission.
5. When the journal table status is pending, you must provide the CloudTrail ARN to connect the journal table to the created infrastructure. Follow the instructions to deploy the CloudFormation stack using the AWS CLI or AWS Management Console and then return to the Workload Factory console to enter the CloudTrail ARN.

Delete the S3 access points for a volume in NetApp Workload Factory

Delete existing S3 access points from a volume in NetApp Workload Factory.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > Volumes explorer**.
3. Select the actions menu for the volume you want to manage S3 access points for, then select **Advanced actions**, and then **Manage S3 access points**.
4. Select the actions menu for the S3 access point to delete and then select **Delete**.
5. In the **Delete an S3 access point** dialog, select **Delete** to delete the S3 access point from the volume.

Workload Factory doesn't delete the inventory or journal table, but you can remove them manually from the AWS Management Console.

Related information

- [Set up the journal table infrastructure](#)

Set up the Journal table infrastructure for NetApp Workload Factory

Set up the Journal table infrastructure to capture and store audit logs of user access

events and object operations across Amazon FSx for ONTAP volume access points. Several steps are necessary to set up the infrastructure for AWS services like AWS CloudTrail, AWS CloudWatch, AWS S3 Buckets, AWS CloudWatch log group, AWS Identity and Access Management (IAM), and AWS S3 Tables so that the log events travel through the pipeline correctly and are read by Workload Factory.

About this task

The Journal table feature captures S3 data-plane events (PutObject, GetObject, DeleteObject, etc.) for monitored FSx for ONTAP S3 access points. It uses a chain of AWS services deployed into your AWS account. When you set up the infrastructure correctly, it connects to the FSx for ONTAP volume access point and establishes the pipeline that captures user access and object operation audit events in the Journal table.

The following table lists the AWS services that are part of the infrastructure, their respective resource name patterns, and the purpose of the service in the pipeline.

AWS service	Resource name pattern	Purpose
AWS CloudFormation	netapp-metadata-*	Deploys all infrastructure as a stack
AWS S3 bucket	netapp-metadata-cloudtrail-events-logs-{uuid}	Stores raw CloudTrail log files
AWS CloudTrail	netapp-metadata-journal-data-events-trail-{uuid}	Captures S3 data events for specific access points
AWS CloudWatch log group	netapp-metadata-journal-data-events-{uuid}	Receives CloudTrail events as structured log entries
IAM roles	- netapp-metadata-cloudtrail-cw-role-{uuid} - netapp-metadata-s3table-integration-role-{uuid}	Active integration
ObservabilityAdmin	S3TableIntegration	Bridges CloudWatch Logs into an S3 Tables table
S3 Tables	aws-cloudwatch bucket → logs.aws_cloudtrail__data	Stores structured, queryable CloudTrail events in Iceberg format

The {uuid} is a random 8-character identifier generated when the template is created.

Before you begin

To enable the Journal table feature, complete these steps:

- Have an existing volume with an S3 access point. [Create a volume with an S3 access point](#)
- Set network configuration to **Internet** for the S3 access point. [Edit network configuration for the S3 access point](#).
- [Grant the operations and remediation permissions](#) to your Workload Factory credentials.
- Add the following IAM policy permissions to the AWS account you use to run the CloudFormation deployment to set up the Journal table.

▼ IAM policy permissions for Journal table setup

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "CFNStack",
      "Effect": "Allow",
      "Action": [
        "cloudformation:CreateStack",
        "cloudformation:DescribeStacks",
        "cloudformation:DescribeStackEvents"
      ],
      "Resource": "arn:aws:cloudformation:*:*:stack/netapp-
metadata-*/*"
    },
    {
      "Sid": "StarResources",
      "Effect": "Allow",
      "Action": [
        "cloudformation:GetTemplateSummary",
        "cloudtrail:DescribeTrails",
        "logs:DescribeLogGroups",
        "logs:ListSourcesForS3TableIntegration",
        "observabilityadmin:CreateS3TableIntegration",
        "observabilityadmin:GetS3TableIntegration",
        "observabilityadmin:TagResource",
        "observabilityadmin:ListTagsForResource"
      ],
      "Resource": "*"
    },
    {
      "Sid": "S3Bucket",
      "Effect": "Allow",
      "Action": [
        "s3:CreateBucket",
        "s3:PutBucketPolicy",
        "s3:PutBucketTagging"
      ],
      "Resource": "arn:aws:s3:::netapp-metadata-*"
    },
    {
      "Sid": "IAMRoles",
      "Effect": "Allow",
      "Action": [
        "iam:CreateRole",
        "iam:PutRolePolicy",

```

```

        "iam:TagRole",
        "iam:GetRole"
    ],
    "Resource": "arn:aws:iam::*:role/netapp-metadata-*"
},
{
    "Sid": "PassRole",
    "Effect": "Allow",
    "Action": "iam:PassRole",
    "Resource": "arn:aws:iam::*:role/netapp-metadata-*",
    "Condition": {
        "StringEquals": {
            "iam:PassedToService": [
                "cloudtrail.amazonaws.com",
                "logs.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "CloudTrail",
    "Effect": "Allow",
    "Action": [
        "cloudtrail:CreateTrail",
        "cloudtrail:StartLogging",
        "cloudtrail:AddTags",
        "cloudtrail:PutEventSelectors"
    ],
    "Resource": "arn:aws:cloudtrail::*:trail/netapp-metadata-
* "
},
{
    "Sid": "CWLogGroup",
    "Effect": "Allow",
    "Action": [
        "logs:CreateLogGroup",
        "logs>DeleteLogGroup",
        "logs:PutRetentionPolicy",
        "logs:TagResource",
        "logs:AssociateSourceToS3TableIntegration"
    ],
    "Resource": "arn:aws:logs::*:log-group/netapp-metadata-*"
},
{
    "Sid": "S3Table",
    "Effect": "Allow",

```

```

        "Action": [
            "s3tables:CreateTableBucket",
            "s3tables:PutTableBucketEncryption",
            "s3tables:PutTableBucketPolicy"
        ],
        "Resource": "arn:aws:s3tables:*:*:bucket/aws-cloudwatch"
    }
}

```

Set up the journal table infrastructure

Set up the infrastructure to capture AWS service events from the S3 access point in the journal table.

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 with the volume to update, then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. From the Volumes tab, select the actions menu for the volume you want to manage S3 access points for, then select **Advanced actions**, and then **Manage S3 access points**.
6. From the **Manage S3 access points** screen, select the actions menu and then select **Edit access point**.
7. In the **Edit S3 access point** dialog, ensure that the network configuration is set to **Internet**.
8. Follow the instructions in the dialog to set up the infrastructure for the journal table feature.
9. Download the CloudFormation template.
10. Deploy the CloudFormation stack in your AWS account.
 - a. Save the template JSON to a file.
 - b. Deploy the template using the AWS CLI or AWS Management Console.
 - c. Wait for the stack to reach `CREATE_COMPLETE` status.
 - d. Get the CloudTrail ARN from the stack outputs.
11. Return to the Workload Factory console and back to the volume to manage S3 access points for.
12. Select **View details** from the volume actions menu.
13. In the **Journal table** tab, enter the CloudTrail ARN.
14. Select **Apply**.

If any of the steps fail, [troubleshoot the infrastructure setup for the journal table](#) or contact NetApp support for assistance.

Troubleshoot the infrastructure setup for the journal table

You can use the AWS Management Console or the AWS CLI to troubleshoot failures during the

CloudFormation stack deployment and the resources it creates.

After resolving the issue, submit the CloudTrail ARN again to restart journal table setup.

Common CloudFormation deployment failures

Insufficient IAM permissions

The deploying role/user needs a specific set of permissions. Refer to [Before you begin](#) for the full policy.

If the stack fails with an `AccessDenied` or `InsufficientPermission` error, check the stack events and look for the missing permission in the `ResourceStatusReason` field.

CloudTrail limit

AWS enforces a default limit of 5 trails per region. If the account already has 5 trails, the `FsxDataEventTrail` resource will fail with: Maximum number of trails (5) exceeded. You can check the number of Trails on a regional level with the following command:

```
aws cloudtrail describe-trails \
  --no-include-shadow-trails \
  --region <region> \
  --query "length(trailList)"
```

Resolution options

- Option 1: Delete an unused trail in the region to make room.
- Option 2: Use an existing trail. Remove the `FsxDataEventTrail`, `CloudTrailBucket`, `CloudTrailBucketPolicy`, and `CloudTrailToCloudWatchRole` resources from the template before deploying the CloudFormation stack again. Then pass the ARN of your existing trail during the initiation step. The existing trail must have a CloudWatch Log Group configured, an `S3TableIntegration` associated with the log group, and be logging data events.

S3 Table integration already exists

If the account already has an `S3TableIntegration` for the `aws_cloudtrail` data source, the `LogsToS3TableIntegration` resource will fail.

Resolution

Remove the `LogsToS3TableIntegration` and `S3TableIntegrationRole` resources from the template before deploying the CloudFormation stack again. The system automatically uses the existing integration as long as you configure it for `aws_cloudtrail` data events.

To check for an existing integration:

```
aws observabilityadmin list-s3-table-integrations --region <your-region>
```

S3 bucket name already exists

The bucket name `netapp-metadata-cloudtrail-events-logs-{uuid}` is globally unique. If it collides,

re-request the template to get a new UUID.

IAM role already exists

If a previous partial deployment left behind IAM roles with the `netapp-metadata-*` name pattern, the stack will fail on role creation. Delete the orphaned roles first:

```
aws iam delete-role-policy \
  --role-name netapp-metadata-cloudtrail-cw-role-<uuid> \
  --policy-name <policy-name>
aws iam delete-role \
  --role-name netapp-metadata-cloudtrail-cw-role-<uuid>
```

Failures after Journal table enablement

After submitting the CloudTrail ARN, Workload Factory validates the entire resource pipeline by sending a *seed*, or test, event automatically. If successful, the seed event arrives in the S3 Tables table. The test takes approximately 10 minutes.

If the test validation fails, you might get one of the following error messages:

Error message	Meaning
Table <code>aws_cloudtrail__data</code> was not created in {bucket}. Verify s3table permissions.	S3TableIntegration did not create the CloudWatch-managed table. The pipeline between CloudWatch Logs and S3 Tables is broken.
Table exists, but the journal seed event does not appear. Verify CloudTrail and CloudWatch permissions.	The table exists but the specific seed event never arrived. Pipeline is broken between CloudTrail and the S3 Tables table.
Failed to initiate journal setup. ...	An error occurred during the background seed/poll flow. Check the trailing message for details.

Resolution steps

When the journal reaches FAILED, trace the seed event forward through the pipeline stages to identify exactly where it stopped. Each step maps to a specific AWS resource created by the template.

1. Check the CloudTrail S3 Bucket.

The trail writes raw event logs to the S3 bucket `netapp-metadata-cloudtrail-events-logs`. Look for recent log files.

If no log files exist, then CloudTrail is not capturing events. Check the following:

- The trail is logging (`IsLogging: true`)
- The advanced event selectors include the correct access point ARN
- The advanced event selectors include the filters for `eventCategory = Data` and `resources.type = AWS::S3::AccessPoint`

2. Check the CloudWatch Log Group.

The trail also delivers events to the CloudWatch Log Group. The log group name starts with `netapp-metadata-journal-data-events-<uuid>`.

- If the log group is empty, then CloudTrail is not delivering events to CloudWatch. Check that the `CloudTrailToCloudWatchRole` IAM role exists and has `logs:CreateLogStream` and `logs:PutLogEvents` permissions, and that the trail is configured with the correct `CloudWatchLogsLogGroupArn` and `CloudWatchLogsRoleArn`.
- If the seed event appears in the log group, then the problem is downstream — proceed to Step 3.

3. Check the S3 Tables table (`aws-cloudwatch`).

The `S3TableIntegration` automatically creates a table bucket called `aws-cloudwatch` and populates a table at `logs.aws_cloudtrail__data`. This table is only created after the first event flows through.

- If the `aws-cloudwatch` table bucket does not exist, then the `S3TableIntegrationRole` is missing permissions. It needs `s3tables:CreateTableBucket`, `s3tables:PutTableBucketEncryption`, and `s3tables:PutTableBucketPolicy` — all scoped to `arn:aws:s3tables:*:*:bucket/aws-cloudwatch`.
- If the table bucket exists but `logs.aws_cloudtrail__data` does not, then the integration is not routing events. The integration must show `Status: ACTIVE` and include `aws_cloudtrail` as a log source.
- If the table exists but the seed event is not in it, then the event may still be in transit. S3 Tables ingestion has some latency. Wait a few more minutes. If it still does not appear after 15-20 minutes, the integration may be broken.

4. Query the seed event directly.

- Open the S3 Tables in the AWS Management Console.
- Navigate to the `aws-cloudwatch` table bucket → `aws_cloudtrail__data` table, and use the **Preview** button to run a quick query directly in the browser.
- If the event is present in the table but the journal still shows `FAILED`, then the polling window may have expired before the event arrived.

After resolving the issue, return to the Workload Factory console. Retry [initiating the journal table setup](#) by submitting the Trail ARN again.

- If setup continues to fail, contact NetApp support for assistance.

Permissions reference for the journal table setup

The IAM role that deploys the CloudFormation stack to enable the Journal table feature needs the following permissions. Refer to [Before you begin](#) for a copiable JSON policy with the required permissions.

Stack Operations

Permission	Resource	Why
<code>cloudformation:CreateStack</code>	<code>arn:aws:cloudformation:*:*:stack/netapp-metadata-/</code>	Create the stack
<code>cloudformation:DescribeStacks</code>	<code>arn:aws:cloudformation:*:*:stack/netapp-metadata-/</code>	Monitor stack status

Permission	Resource	Why
cloudformation:DescribeStackEvents	arn:aws:cloudformation:*:*:stack/netapp-metadata-/*	Diagnose resource-level failures cloudformation:GetTemplateSummary * Pre-flight template validation

CloudTrail

Permission	Resource	Why
cloudtrail:CreateTrail	arn:aws:cloudtrail:*:*:trail/netapp-metadata-*	Create the trail
cloudtrail:StartLogging	arn:aws:cloudtrail:*:*:trail/netapp-metadata-*	Enable logging
cloudtrail:AddTags	arn:aws:cloudtrail:*:*:trail/netapp-metadata-*	Apply identification tag
cloudtrail:PutEventSelectors	arn:aws:cloudtrail:*:*:trail/netapp-metadata-*	Configure data event capture
cloudtrail:DescribeTrails	*	Resolve trail ARN for stack output

S3

Permission	Resource	Why
s3:CreateBucket	arn:aws:s3:::netapp-metadata-*	Create the CloudTrail log bucket
s3:PutBucketPolicy	arn:aws:s3:::netapp-metadata-*	Allow CloudTrail to write logs
s3:PutBucketTagging	arn:aws:s3:::netapp-metadata-*	Apply identification tag

IAM

Permission	Resource	Why
iam:CreateRole	arn:aws:iam:*:*:role/netapp-metadata-*	Create both IAM roles
iam:PutRolePolicy	arn:aws:iam:*:*:role/netapp-metadata-*	Attach inline policies
iam:TagRole	arn:aws:iam:*:*:role/netapp-metadata-*	Apply identification tag
iam:GetRole	arn:aws:iam:*:*:role/netapp-metadata-*	Confirm role is active
iam:PassRole	arn:aws:iam:*:*:role/netapp-metadata-* (condition: PassedToService = cloudtrail.amazonaws.com, logs.amazonaws.com)	Pass roles to CloudTrail and CloudWatch Logs

CloudWatch Logs

Permission	Resource	Why
logs:CreateLogGroup	arn:aws:logs:*:*:log-group:netapp-metadata-*	Create the log group
logs>DeleteLogGroup	arn:aws:logs:*:*:log-group:netapp-metadata-*	Clean up log group if creation failed
logs:PutRetentionPolicy	arn:aws:logs:*:*:log-group:netapp-metadata-*	Set 30-day retention
logs:TagResource	arn:aws:logs:*:*:log-group:netapp-metadata-*	Apply identification tag
logs:AssociateSourceToS3TableIntegration	arn:aws:logs:*:*:log-group:netapp-metadata-*	Link CloudTrail source to S3 Tables
logs:DescribeLogGroups	*	Check log group existence
logs:ListSourcesForS3TableIntegration	*	Confirm integration association

ObservabilityAdmin

Permission	Resource	Why
observabilityadmin:CreateS3TableIntegration	*	Create the CloudWatch → S3 Tables bridge
observabilityadmin:GetS3TableIntegration	*	Confirm integration is active
observabilityadmin:TagResource	*	Apply identification tag
observabilityadmin:ListTagsForResource	*	Drift detection

S3 Tables

Permission	Resource	Why
s3tables:CreateTableBucket	arn:aws:s3tables:*:*:bucket/aws-cloudwatch	Create the S3 Tables bucket (via integration role)
s3tables:PutTableBucketEncryption	arn:aws:s3tables:*:*:bucket/aws-cloudwatch	Set AES256 encryption
s3tables:PutTableBucketPolicy	arn:aws:s3tables:*:*:bucket/aws-cloudwatch	Allow CloudWatch Logs access

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.