



Get started

NetApp Console local deployment

NetApp
August 10, 2026

This PDF was generated from <https://docs.netapp.com/us-en/console-local/concept-netapp-console-local.html> on August 10, 2026. Always check docs.netapp.com for the latest.

Table of Contents

- Get started 1
 - Learn about NetApp Console local deployment 1
 - Deploy and operate Console in your own infrastructure 1
 - Automate storage operations with APIs and service accounts 1
 - Control access with RBAC and Active Directory integration 1
 - Organize fleets and delegate storage administration 1
 - Track and export administrative activity for compliance 2
 - Provision storage consistently with storage class policies 2
 - Monitor storage class compliance and remediate drift automatically 2
 - Monitor storage health and receive alerts across fleets 2
 - Provision storage and investigate alerts with natural language 2
 - Back up and restore storage with NetApp Backup and Recovery 3
 - Learn about NetApp Console local deployment identity and access management 3
 - What IAM means in NetApp Console local deployment 3
 - How authentication works 3
 - How authorization works 4
 - How the IAM hierarchy works 4
 - Member security and credentials 4
 - Audit IAM activity 5
 - Next steps with IAM in NetApp Console 5
 - Learn about fleet management in NetApp Console local deployment 5
 - Why fleet management matters 5
 - How fleets organize your resources 6
 - Common fleet organization patterns 6
 - How fleet management enables access control 6
 - How storage management works 7
 - Provisioning approaches 7
 - Where to go next 7
 - Learn about storage classes and policies in NetApp Console local deployment 8
 - What storage class policies are 8
 - What storage classes are 8
 - How storage classes enforce standards 8
 - Storage class drift and compliance 9
 - End-to-end workflow 9
 - How storage classes work with fleets 10
 - Where to go next 10
 - Learn about LLM integration in NetApp Console local deployment 10
 - What you can do with AI-assisted storage management 10
 - Choose read-only or read/write access for the Console assistant 11
 - Understand what controls Console assistant actions 11
 - Choose an LLM provider path 11
 - Understand where LLM requests go 11
 - Protect LLM credentials and control admin access 11

Connect your LLM	11
Learn about NetApp Backup and Recovery in NetApp Console local deployment	12

Get started

Learn about NetApp Console local deployment

NetApp Console local deployment lets you host and run the NetApp Console autonomous control plane in your own infrastructure.

You get unified storage management, policy enforcement, fleet-scale automation, and AI-assisted operations. No data, metadata, or management traffic leaves your environment.

Deploy and operate Console in your own infrastructure

NetApp Console local deployment runs in your own infrastructure, so your team controls deployment, connectivity, and daily operations. You deploy the Console agent, maintain the Console virtual machine, and manage the network paths required for monitoring and management traffic. This gives enterprise administrators full control of operational boundaries while keeping management data inside the environment.

- [Install NetApp Console local deployment using an OVA in vCenter.](#)
- [Maintain your vCenter or ESXi host for NetApp Console local deployment.](#)
- [Learn about ports for the on-premises Console agent in NetApp Console local deployment.](#)

Automate storage operations with APIs and service accounts

NetApp Console local deployment supports API-based integrations so your organization can automate repeatable storage operations. Service accounts let applications authenticate and operate with assigned roles. The same role-based access control and audit controls that apply to interactive users govern those actions. This supports enterprise automation while keeping access scoped and accountable.

- [Add members to your NetApp Console local deployment organization.](#)
- [Learn about the API for NetApp Console IAM](#)

Control access with RBAC and Active Directory integration

NetApp Console local deployment provides zero-trust role-based access control (RBAC) that limits what users can see and do within the fleets and resources they manage. You can integrate with Active Directory so users sign in with their existing corporate credentials, and local users can add multi-factor authentication (MFA) for another layer of verification. Access governance stays aligned with your organization's broader identity and access management practices.

- [Learn about identity and access management in NetApp Console local deployment.](#)
- [Learn about secure access to NetApp Console local deployment.](#)

Organize fleets and delegate storage administration

NetApp Console local deployment scales administration by organizing resources into folders and fleets. You can map fleets to environments, business units, or regions, then delegate administration at organization, folder, or fleet scope to keep ownership clear. This structure helps large storage teams distribute work without losing policy consistency or governance.

- [Learn about folders and fleets in NetApp Console local deployment.](#)

- [Learn about storage fleets in NetApp Console local deployment.](#)

Track and export administrative activity for compliance

Every administrative action generates an audit record. Security and compliance teams can use these records to investigate incidents, show control effectiveness, and satisfy audit requirements. Export audit logs to your syslog server through a webhook for centralized monitoring, longer retention, and analysis. Because NetApp Console local deployment runs in your own environment, log data never leaves your infrastructure.

- [Audit NetApp Console local deployment activity.](#)

Provision storage consistently with storage class policies

NetApp Console local deployment enforces consistent storage behavior through storage classes—templates that bundle performance, capacity, and tiering policies into reusable definitions. Administrators define storage standards once. Administrators and automated workflows use those approved classes for all provisioning activity across your environment. This stops configuration drift, reduces the time junior administrators spend on provisioning decisions, and guarantees that every workload meets your organization's standards immediately. After provisioning, Console local deployment continues to monitor each workload for compliance.

- [Learn about managing storage with NetApp Console local deployment.](#)

Monitor storage class compliance and remediate drift automatically

NetApp Console local deployment monitors every workload against the storage class used to provision it. When a workload drifts—because of a direct cluster configuration change or degrading performance—NetApp Console local deployment flags the violation immediately. Administrators can follow guided remediation steps or configure automated remediation to resolve common drift conditions without manual intervention. This continuous enforcement reduces audit risk and keeps your environment compliant between scheduled reviews.

Monitor storage health and receive alerts across fleets

NetApp Console local deployment gives you a single place to watch the health and performance of every cluster you manage. Fleet-wide dashboards display clusters and volumes that need attention. Custom dashboards let administrators build monitoring views that match their responsibilities. The Workload Analyzer correlates latency, throughput, resource utilization, and configuration changes over time to help you pinpoint root causes before issues affect workloads.

Configure email and webhook delivery channels so alerts reach the right teams when conditions change. Organization admins set up destinations, and storage admins apply them in alert rules so response paths match your operational model. This helps enterprise teams respond faster to capacity, performance, and protection events.

- [Learn about monitoring storage health in NetApp Console local deployment.](#)
- [Set up notification delivery channels in NetApp Console local deployment.](#)
- [Configure notification rules for alerts in NetApp Console local deployment.](#)

Provision storage and investigate alerts with natural language

NetApp Console local deployment can connect to your own large language model (LLM) to provide AI-assisted storage management. You can describe a workload in plain language to get a provisioning plan, ask the Console to investigate an active alert, or get contextual answers about your storage environment. Every AI action stays within the guardrails of your storage classes and the role-based access controls that apply to your

account, and you must confirm sensitive operations before they proceed. Console local deployment sends requests only to the LLM endpoint your administrators configure.

- [Learn about LLM integration in NetApp Console local deployment.](#)

Back up and restore storage with NetApp Backup and Recovery

Beyond provisioning and monitoring, NetApp Console local deployment gives you access to NetApp Backup and Recovery so you can protect your data from the same interface. You can back up and restore your data to support your protection and recovery requirements. Managing data protection alongside your storage keeps governance consistent and reduces the number of tools your teams need to operate.

- [Learn about data services in NetApp Console local deployment.](#)

Learn about NetApp Console local deployment identity and access management

Identity and access management (IAM) in NetApp Console local deployment controls who can sign in, how identities are verified, what resources users can access, and what actions they can perform. In NetApp Console local deployment, IAM includes role-based access control (RBAC), multi-factor authentication (MFA), and directory-backed authentication, as well as the hierarchy and audit model that support delegated administration.

Use this page to understand the NetApp Console local deployment IAM model at a high level. For detailed hierarchy planning examples, [Learn about folders and fleets in NetApp Console local deployment.](#)



You must have the *Super admin*, *Organization admin*, or *Folder or fleet admin* roles to manage IAM in NetApp Console.

What IAM means in NetApp Console local deployment

NetApp Console local deployment IAM combines identity verification, access control, delegation, and auditability:

- *Authentication* determines how users sign in, whether with local credentials or through your directory configuration.
- *Authorization* determines what members can do after they sign in, based on predefined roles and the scope where you assign them.
- *Hierarchy and scoping* determine which folders, fleets, and resources a member can access.
- *Member and credential management* determine how you add users, service accounts, and how you manage MFA and service account secrets.
- *Audit and compliance tracking* records IAM-related activity so you can review who changed access and when.

How authentication works

NetApp Console local deployment supports both local identities and external identity integration.

You can integrate NetApp Console local deployment with Active Directory so users sign in with their existing corporate credentials. This eliminates the need for separate passwords in Console local deployment and lets administrators look up directory users when assigning access.

For local users, MFA adds another verification step to reduce the risk of unauthorized access if credentials are compromised.

To learn more about authentication and secure sign-in options, [Learn about secure access in NetApp Console local deployment](#).

How authorization works

After a user signs in, NetApp Console local deployment uses RBAC to determine what that user can see and do.

Active Directory or LDAP integration handles authentication. NetApp Console local deployment authorization remains separate: administrators assign predefined roles at the organization, folder, or fleet level to control what each member can access.

The same role grants different effective access depending on the scope where you assign it. This model lets you give broad access to central administrators while limiting regional or team-level administrators to the fleets they manage.

Roles that you assign at the organization or folder level are inherited by child scopes. This inheritance model is a key part of how NetApp Console delegates access across folders and fleets.

NetApp designs NetApp Console local deployment roles with least-privilege principles so each role includes only the permissions needed for its tasks.

[Learn about role-based access control and role inheritance in NetApp Console local deployment](#).

How the IAM hierarchy works

NetApp Console local deployment uses a hierarchy to group resources and scope access. The organization is at the top, then folders, then fleets, and then the resources (including possible sub-folders) associated with those fleets.

This hierarchy is what makes delegation possible. For example, an Organization admin can manage access across the entire organization, while a Folder or fleet admin can manage only the resources and members within the part of the hierarchy they own.

NetApp Console IAM is built on three types of components: organizational components that define the hierarchy, resources that are assigned within that hierarchy, and members and roles that control who can access what.

Because roles inherit through this hierarchy, your folder and fleet design directly affects how broadly each access assignment applies.

[Learn how to add fleets to your organization..](#)

Member security and credentials

IAM in NetApp Console local deployment also includes operational controls for securing member access after accounts exist.

For local users, administrators can help users recover access by directing password reset, removing or temporarily disabling MFA, and reviewing local-user MFA behavior. For service accounts, administrators can recreate credentials when secrets are lost or rotated.

These controls are part of IAM because they determine how identities remain secure over time, not just how access is assigned initially.

[Learn how to manage member security.](#)

Audit IAM activity

IAM in NetApp Console local deployment includes the ability to review and export access-related activity.

From the Audit page, you can review actions related to managing your organization, such as adding members, creating fleets, and associating resources. You can also filter audit records by the Tenancy service to focus on IAM-related changes, or export audit logs to an external system.

Auditability is an important IAM principle because it lets you verify who changed access, when the change happened, and whether the change was successful.

[Learn how to audit NetApp Console local deployment activity.](#)

Next steps with IAM in NetApp Console

- [Get started with identity and access in NetApp Console](#)
- [Learn about secure access in NetApp Console local deployment](#)
- [Learn about RBAC in NetApp Console local deployment](#)
- [Monitor or audit Console local deployment activity](#)
- [Learn about the API for NetApp Console IAM](#)

Learn about fleet management in NetApp Console local deployment

Fleet management in NetApp Console local deployment is the practice of organizing storage systems into logical groups so you can apply consistent policies, control access, and monitor health across related clusters. Instead of managing each cluster independently, fleet management lets you treat your fleet as a common pool of resources to support your data storage goals.

As your storage environment grows, managing individual clusters becomes impractical. Fleet management solves this by giving you a structured way to group related systems, delegate responsibilities, and ensure every cluster operates under the same standards.

Why fleet management matters

Fleet management addresses three core challenges:

- **Simplification through abstraction** - Fleet management abstracts away the complexity of managing individual storage infrastructure. Instead of dealing with cluster-by-cluster configuration, you manage your storage estate as a unified whole, reducing operational overhead and decision fatigue.

- **Consistency and scalability** - Without clear organization, different teams make different decisions for similar workloads, leading to fragmented configurations. Fleet management lets you apply standards across dozens of systems at once, ensuring new workloads start from the same baseline across teams and fleets.
- **Governance and control** - As teams grow, you need clear boundaries for who can manage what. Fleet management provides those boundaries through organizational structure and access control, ensuring that storage decisions remain governed and auditable.

How fleets organize your resources

Your organization's resource hierarchy consists of folders and fleets:

For a detailed overview of the hierarchy and access model, see [Learn about folders and fleets in NetApp Console local deployment](#).

- **Organization** - The top level representing your company.
- **Folders** - Help you organize related fleets. For example, you might create a folder for each region or business unit, then create fleets within each folder. Folders can contain other folders or fleets. When you assign a role at the folder level, members inherit that role for all fleets within the folder.
- **Fleets** - Group the storage systems you manage together. You associate storage systems with fleets and assign members to fleets to grant them access.



Folders are an organizational tool and are not visible to members who do not have IAM permissions such as Organization admin, Folder admin, or Fleet admin roles. Members access fleets, not folders.

Common fleet organization patterns

How you organize fleets depends on your operational model:

- **By environment** - Create separate fleets for production, development, and test workloads. This keeps production storage under tighter policies while allowing more flexibility in lower environments.
- **By business unit** - Group clusters that serve a specific department, such as finance, engineering, or HR, into a dedicated fleet. You can then assign storage management responsibilities to team members within that business unit without exposing other fleets to them.
- **By location or data center** - When clusters are distributed across sites, organizing by location lets you monitor site-level health, apply region-specific policies, and track capacity consumption per site.
- **By application tier** - Group clusters that host a specific class of workload, such as databases, file shares, or virtual machines, so you can apply consistent standards tuned to that workload type across the entire fleet.



Use folders to add another level of organization. For example, create a folder for each region and then create environment-based fleets within each regional folder.

How fleet management enables access control

Fleets and folders serve as boundaries for user access and administrative responsibility:

- **Folder-level access** - When you assign a role at the folder level, the member inherits that role for all fleets and resources within the folder. This lets you delegate administrative responsibilities without granting

access to the entire organization.

- **Fleet-level access** - When you assign a role at the fleet level, the member has access only to the storage systems within that fleet. This lets you delegate storage management to team members or application owners without exposing unrelated parts of your infrastructure.

Console local deployment provides fleet-level health and performance monitoring so you can see the status of all clusters in a fleet from a single view, identify issues early, and act before they affect workloads.

How storage management works

Storage management is the practice of defining storage standards, applying them consistently, and operating workloads so they stay aligned over time. In Console local deployment, those standards are expressed as storage class policies and storage classes, scoped to fleets, and enforced through provisioning workflows governed by RBAC and audit logging.

Console local deployment connects four concepts into one workflow:

- **Fleets** define the scope where you manage storage. A fleet groups systems so you can control access, apply standards consistently, and manage resources as a unit.
- **Storage class policies** define standards as reusable building blocks (performance, capacity, security, data protection).
- **Storage classes** express workload intent. A storage class is a named template assembled from one or more policies.
- **Provisioning workflows** create storage within guardrails. Different workflows make configuration decisions differently, but all can enforce storage classes and remain governed by RBAC and audit logging.

Provisioning approaches

Console local deployment supports three provisioning approaches, all governed by RBAC, audit logging, and storage class standards:

- **Manual provisioning** - You select the target system and specify configuration details directly. Use this when you need explicit control over system selection and configuration.
- **Automated provisioning** - You provide workload inputs and optionally select a storage class. Console local deployment recommends best-fit placement and applies class-driven settings to reduce manual decisions.
- **AI-assisted (agentic) provisioning** - You describe what you need in natural language. Console local deployment proposes a plan constrained by RBAC and storage classes, then provisions only after you review and approve.

Where to go next

- To understand storage standards, see [Learn about storage classes and policies in NetApp Console local deployment](#).
- To create and manage fleets, see [Manage folders and fleets](#).
- [Provision storage manually](#)
- [Provision storage automatically](#)
- [Provision storage with AI assistance](#)

Learn about storage classes and policies in NetApp Console local deployment

A storage class is a reusable template that defines how storage should behave. When you provision storage using a storage class, NetApp Console local deployment enforces the standards encoded in that class automatically without requiring administrators to make individual configuration decisions for every workload.

Storage classes are built from storage class policies, which define individual dimensions of storage behavior. Together, they let you capture your organization's storage standards once and apply them consistently across fleets.

What storage class policies are

A storage class policy defines one dimension of storage behavior. Policies are reusable building blocks that you combine to create storage classes.

Common policy types include:

- **Performance policies** - Define latency targets, IOPS, or throughput requirements.
- **Capacity policies** - Define provisioning mode, threshold alerts, or growth limits.
- **Security policies** - Define encryption, compliance, or access control requirements.
- **Data protection policies** - Define snapshot schedules, replication targets, or retention periods.

You define policies independently, then combine them when you build a storage class. This lets you reuse the same performance policy across multiple classes, or update a capacity policy in one place and apply that change to every class that uses it.

What storage classes are

A storage class is a named template assembled from one or more policies. It captures your organization's storage standards for a specific type of workload.

For example, you might create a **Production Database** storage class that combines high-performance, high-availability, and strict data protection policies, or a **Development File Share** storage class that combines moderate performance, flexible capacity, and basic data protection policies.

Storage administrators define storage classes to encode enterprise requirements. All provisioning activity, whether done manually, through guided workflows, or with automation, draws from the library of classes those administrators have approved.

How storage classes enforce standards

When you provision storage, you select a storage class rather than configuring individual settings. Console local deployment uses the policies in that class to identify which clusters in your fleet have the capacity and performance headroom to support the workload, recommend the best placement option, and apply class-driven settings automatically at provisioning time.

After provisioning, Console local deployment continuously evaluates each workload against its storage class standards.

Storage class drift and compliance

When a workload no longer matches its storage class intent, Console local deployment flags it for investigation and remediation.

Issues fall into two categories:

- **Configuration drift:** Storage settings governed by the storage class policy no longer match the intended configuration. This can occur when changes are made directly on the ONTAP cluster or outside standard provisioning workflows.
- **Performance non-compliance:** The workload's runtime behavior no longer meets the storage class performance intent (for example, sustained pressure near a QoS limit or elevated tail latency).

An analysis view explains what changed and why. Guided remediation actions (for example, **Fix it**) may be available to help restore compliance. Some remediations can be applied in place, while others may require aligning the workload to a different storage class to restore the intended behavior.

Where to investigate

You can typically investigate drift and non-compliance from one of these locations (availability depends on your deployment and permissions):

- **Storage classes:** Drift / Compliance
- **Alerts:** Analyze

For step-by-step instructions, see [Remediate storage class drift](#).

End-to-end workflow

The following steps describe the process for using storage classes to standardize and govern storage in your environment:

1. **Create storage class policies** - Policies define the individual dimensions of storage behavior (performance targets, capacity settings, security requirements, data protection schedules). Create policies before building a storage class.
2. **Create a storage class** - Assemble a storage class by combining one policy from each applicable category. You can use a predefined storage class or create a custom one for your organization's standards.
3. **Associate the storage class with a fleet** - After you create a storage class, associate it with the fleet where it will be used for provisioning and compliance monitoring.
4. **Provision storage using the storage class** - When provisioning a volume or LUN, select a storage class instead of configuring individual settings. Console local deployment uses the class to recommend the best-fit cluster placement, then provisions with the settings defined in the class.
5. **Monitor workloads for drift** - Console local deployment continuously evaluates each workload against the standards encoded in its storage class. If configuration drift or performance non-compliance occurs, it flags the issue and may raise an alert.
6. **Remediate drift to restore compliance** - When drift or performance non-compliance is detected, you can follow guided steps to correct the issue manually, let Console local deployment resolve common drift conditions automatically, or disassociate a workload from its storage class if you need to change its settings.

How storage classes work with fleets

Storage classes are associated with fleets. When you associate a storage class with a fleet, that class becomes available for all provisioning within the fleet. This ensures that every workload provisioned in the fleet follows the same standards, making fleets the primary way to enforce consistent storage behavior across related clusters.

For details on how to organize fleets, see [Learn about fleet management in NetApp Console local deployment](#).

Where to go next

- To understand fleet organization, see [Learn about fleet management in NetApp Console local deployment](#).
- To create policies and storage classes, see [Manage storage classes and policies](#).
- [Provision storage manually](#)
- [Provision storage automatically](#)
- [Provision storage with AI assistance](#)
- To analyze and remediate drift, see [Remediate storage class drift](#)

Learn about LLM integration in NetApp Console local deployment

NetApp Console local deployment connects to a large language model (LLM) for AI-assisted storage management. After setup, users can use natural language to provision storage, investigate alerts, and get storage guidance.



This documentation regarding connecting an LLM to the Console assistant to enable AI features is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

AI-assisted management lets users describe requests in plain language while Console local deployment applies your storage policies.

Console local deployment sends LLM requests only to the endpoint you configure.

What you can do with AI-assisted storage management

When you enable LLM integration, Console local deployment provides three capabilities through the Console assistant:

- **Storage provisioning** Describe workload requirements in plain language. Console local deployment recommends and runs a provisioning plan based on your storage classes.
- **Alert response** Ask Console local deployment to investigate an active alert. It analyzes the issue and suggests or runs an action based on assigned permissions.
- **Search and guidance** Ask questions about your storage environment or ONTAP administration. Console local deployment returns contextual answers from documentation and current fleet state.

Choose read-only or read/write access for the Console assistant

Users choose an assistant access mode based on the outcome they want:

- **Read only** mode for guidance and investigation without making infrastructure changes.
- **Read/write** mode for guided execution. Console local deployment shows the proposed action, requests confirmation, and then runs the change.

Understand what controls Console assistant actions

Console local deployment controls Console assistant actions through the same governance model used across the platform:

- Role-based access controls limit what each user can see and do.
- Storage policies constrain provisioning and related actions.
- The Console assistant requires confirmation before it runs storage-changing actions.

Choose an LLM provider path

NetApp Console local deployment supports these LLM provider types:

- **OpenAI** for direct OpenAI model access.
- **OpenAI-compatible** for a compatible endpoint, such as an enterprise gateway or proxy service.
- Anthropic for Anthropic's Claude models.

Model availability depends on the endpoint you configure. Select a model from the list in Console local deployment.

For supported model details, see [Learn about supported LLM providers and models in NetApp Console local deployment](#).

Understand where LLM requests go

Data flow depends on the endpoint path you choose:

- If you configure an OpenAI endpoint, Console local deployment sends prompts and context to the OpenAI endpoint.
- If you configure an OpenAI-compatible endpoint, Console local deployment sends prompts and context to the compatible endpoint that your organization configures.

Protect LLM credentials and control admin access

Protect the integration with these controls:

- Treat the API key as a privileged credential and rotate it according to your organization policy.
- Review provider-side usage and alerts to detect unexpected activity.

Connect your LLM

After you make these decisions, continue with [Connect your LLM and enable the NetApp Console assistant](#).

If connection or runtime checks fail, see [Troubleshoot LLM integration](#).

Learn about NetApp Backup and Recovery in NetApp Console local deployment

NetApp Backup and Recovery is a data service that provides efficient, secure, and cost-effective data protection for all your ONTAP workloads, including volumes, databases, virtual machines, and Kubernetes workloads.

A workload is a logical grouping of resources within a system that is managed as a single unit for protection, governance, detection, and recovery. In Backup and Recovery, a workload is the unit you protect. Its meaning depends on the data source: for Kubernetes a workload is an application, for VM environments it's a virtual machine, and for database environments it's a database.

Support for Backup and Recovery is already built in to all ONTAP systems, so there is no need for additional hardware or media gateways. This makes backup operations simple and cost effective. The NetApp Console simplifies the implementation of any backup strategy, including the full spectrum of 3-2-1 backup variants, without needing multiple resource managers or specialized personnel.



NetApp Backup and Recovery is in preview mode for NetApp Console local deployment. The service is not yet generally available, and the features and functionality are subject to change.

[Learn more about NetApp Backup and Recovery..](#)

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.