



Provisioning and managing workloads

Active IQ Unified Manager

NetApp
March 26, 2025

Table of Contents

Provisioning and managing workloads	1
Workloads overview	1
Workloads overview section	2
Data center overview section	2
All Workloads	3
Provisioning workloads	5
Managing Performance Service Levels	6
Guidelines for creating a custom Performance Service Level	11
Creating and editing Performance Service Levels	11
Managing Storage Efficiency Policies	12
Guidelines for creating a custom Storage Efficiency Policy	13
Creating and editing Storage Efficiency Policies	13

Provisioning and managing workloads

The active management feature of Active IQ Unified Manager provides Performance Service Levels, Storage Efficiency Policies, and storage provider APIs for provisioning, monitoring, and managing storage workloads in a data center.



Unified Manager provides this functionality by default. You can disable it from **Storage Management > Feature Settings** if you do not plan to use this functionality.

When enabled, you can provision workloads on the ONTAP clusters managed by your instance of Unified Manager. You can also assign policies, such as Performance Service Levels and Storage Efficiency Policies on the workloads and manage your storage environment based on those policies.

This feature enables the following functions:

- Automatic discovery of storage workloads on the added clusters enabling easy storage workload evaluation and deployment
- Provisioning NAS workloads supporting NFS and CIFS protocols
- Provisioning SAN workloads supporting iSCSI and FCP protocols
- Support for both NFS and CIFS protocols on the same file share
- Management of Performance Service Levels and Storage Efficiency Policies
- Assigning Performance Service Levels and Storage Efficiency Policies to storage workloads

The **Provisioning**, **Storage > Workloads**, and **Policies** options on the left pane of the UI enable you to modify various configurations.

You can perform the following functions by using these options:

- View storage workloads on the **Storage > Workloads** page
- Create storage workloads from the Provision Workload page
- Create and manage Performance Service Levels from Policies
- Create and manage Storage Efficiency Policies from Policies
- Assign policies to storage workloads from the Workloads page

Workloads overview

A workload represents the input/output (I/O) operations of a storage object, such as a volume or LUN. The way the storage is provisioned is based on the expected workload requirements. Workload statistics are tracked by Active IQ Unified Manager only after there is traffic to and from the storage object. For example, the workload IOPS and latency values are available after users start using a database or email application.

The Workloads page displays a summary of the storage workloads of the ONTAP clusters managed by Unified Manager. It provides cumulative at-a-glance information about the storage workloads that conform to the Performance Service Level, as well as the non-conforming storage workloads. It also enables you to assess the total, available, and used capacity and performance (IOPS) of the clusters across your data center.



It is recommended that you assess the number of storage workloads that are non-conforming, unavailable, or not managed by any Performance Service Level, and take the necessary actions to ensure their conformance, capacity usage, and IOPS.

The Workloads page has the following two sections:

- Workloads overview: Provides an overview of the number of storage workloads on the ONTAP clusters managed by Unified Manager.
- Data center overview: Provides an overview of the capacity and IOPS of the storage workloads in the data center. The relevant data is displayed at a data center level and for individual .

Workloads overview section

The workloads overview section provides cumulative at-a-glance information of the storage workloads. The status of the storage workloads is displayed based on assigned and unassigned Performance Service Levels.

- **Assigned:** The following statuses are reported for storage workloads on which Performance Service Levels have been assigned:
 - **Conforming:** Performance of storage workloads is based on the Performance Service Levels assigned to them. If the storage workloads are within the threshold latency defined in the associated Performance Service Levels, they are marked as “conforming”. The conforming workloads are marked in blue.
 - **Non-conforming:** During performance monitoring, storage workloads are marked as “non-conforming” if the storage workloads latency exceeds the threshold latency defined in the associated Performance Service Level. The non-conforming workloads are marked in orange.
 - **Unavailable:** Storage workloads are marked as “unavailable” if they are offline, or if the corresponding cluster is unreachable. The unavailable workloads are marked in red.
- **Unassigned:** Storage workloads that do not have a Performance Service Level assigned to them, are reported as “unassigned”. The number is conveyed by the information icon.

The total workload count is the sum total of the assigned and unassigned workloads.

You can click the total number of workloads displayed in this section, and view them on the Workloads page.

The Conformance by Performance Service Levels subsection displays the total number of available storage workloads:

- Conforming to each type of Performance Service Level
- For which there is a mismatch between the assigned and the recommended Performance Service Levels

Data center overview section

The data center overview section graphically represents the available and used capacity, and IOPS for all of the clusters in the data center. By using this data, you should manage the capacity and IOPS of the storage workloads. The section also displays the following information for the storage workloads across all of the clusters:

- The total, available, and used capacity for all of the clusters in your data center
- The total, available, and used IOPS for all of the clusters in your data center
- The available and used capacity based on each Performance Service Level

- The available and used IOPS based on each Performance Service Level
- The total space and IOPS used by the workloads that have no Performance Service Level assigned

How data center capacity and performance is calculated based on Performance Service Levels

The used capacity and IOPS is retrieved in terms of the total used capacity and performance of all of the storage workloads in the clusters.

The available IOPS is calculated based on the expected latency and recommended Performance Service Levels on the nodes. It includes the available IOPS for all of the Performance Service Levels whose expected latency is less than or equal to their own expected latency.

The available capacity is calculated based on the expected latency and recommended Performance Service Levels on aggregates. It includes the available capacity for all of the Performance Service Levels whose expected latency is less than or equal to their own expected latency.

All Workloads

The All Workloads view displays the list of all the workloads available on the clusters in a data center.

The All Workloads view lists the storage workloads associated with the ONTAP clusters managed by Unified Manager. The page also enables you to assign Storage Efficiency Policies and Performance Service Levels to storage workloads.

When you add clusters to Unified Manager, the storage workloads on each cluster are automatically discovered and displayed on this page, except FlexGroup volumes and its constituents.

Unified Manager begins to analyze the workloads for recommendation (recommended Performance Service Levels) only after I/O operations start on the storage workloads. For the newly-discovered storage workloads on which there have been no I/O operations, the status is “Waiting for I/O”. After I/O operations begin on the storage workloads, Unified Manager begins the analysis and the workload status changes to “Learning...”. After the analysis is complete (within 24 hours from the beginning of the I/O operations), the recommended Performance Service Levels are displayed for the storage workloads.

Using the **Workloads > All Workloads** option, you can perform multiple tasks:

- Add or provision storage workloads
- View and filter the list of workloads
- Assign Performance Service Levels to individual storage workloads
- Assign system-recommended Performance Service Levels to all unassigned workloads
- Assign Storage Efficiency Policies to storage workloads

Adding or provisioning storage workloads

You can add or provision the storage workloads to supported LUNs (supporting both iSCSI and FCP protocols), NFS file shares, and SMB shares.

Viewing and filtering workloads

On the All Workloads screen, you can view all the workloads in your data center or search for specific storage workloads based on either the Performance Service Levels assigned to them, or by workload names. You can

use the filter icon to enter specific conditions for your search. You can search by different filter conditions, such as by the host cluster or storage VM. If you select the **Capacity Total** option, you can search by the total capacity of the workloads. The number of workloads returned might vary depending on the exact filter condition entered, as the capacity is compared at a byte level. The screen displays storage information for the workloads, such as the host cluster and storage VM, along with the Performance Service Level and Storage Efficiency Policy, if any.

The page also enables you to view the performance details of a workload. You can view detailed information about the IOPS, capacity, and latency of the workload by clicking the **Choose / Order Columns** button and selecting specific columns to view. The Performance View column displays the average and peak IOPS for a workload, and you can click the workload analyser icon to view the detailed IOPS analysis. The **Analyze Workload** button on the IOPS Analysis pop-up takes you to the Workload Analysis page, where you can select a time range and view the latency, throughput, and capacity trends for the selected workload. For more information about workload analyzer, see *Troubleshooting workloads using the workload analyzer*

[Troubleshooting workloads using the workload analyzer](#)

Assigning Performance Service Levels to storage workloads

You can assign a Performance Service Level to a single or multiple storage workloads. By using the different navigations on the screen, you can assign a specific or system-recommended Performance Service Level to the workloads.

Assigning system-recommended Performance Service Levels to all unassigned workloads

You can assign Performance Service Levels to all the storage workloads in a data center that have no PSLs assigned, provided there are system-recommended Performance Service Levels available for the workloads. This functionality is not effective for workloads that do not have any system-recommended Performance Service Levels available. To use this functionality, click the **Assign System Recommended PSLs** button. The system internally evaluates the appropriate Performance Service Levels for the unassigned storage workloads and displays the total number of workloads on which the Performance Service Levels can be assigned.

Assigning Storage Efficiency Policies to storage workloads

You can assign a Storage Efficiency Policy to a single or multiple storage workloads. By using the different navigations on the screen, you can assign a specific Storage Efficiency Policy to the storage workloads.

Assigning Performance Service Level and Storage Efficiency Policy together

You can assign a Performance Service Level and Storage Efficiency Policy to a single workload together. Follow these steps:

1. Click the edit icon for the row, and then click **Edit**.

The **Assigned Performance Service Level** and **Storage Efficiency Policy** fields are enabled.

2. Select the required Performance Service Level and Storage Efficiency Policy.
3. Click the check icon to apply the changes.

Analyzing performance and capacity criteria for a workload

You can view performance information about a workload to help with troubleshooting by clicking the bar chart icon in the **Performance View** column. To view performance and capacity charts on the Workload Analysis page to analyze the object, click the **Analyze Workload** button.

Provisioning workloads

The Provision Workload page enables you to create LUNs (supporting both iSCSI and FCP protocols) or file shares (supporting CIFS/SMB and NFS protocols) for an existing cluster and Storage Virtual Machine (storage VM) when there is free space available.

Before you begin

- At least one of SMB, NFS, iSCSI, and FCP services, or all the services should be enabled on your Storage VM. If you are creating LUNs, both iSCSI and FCP must be enabled.
- For selecting and assigning the Performance Service Level and Storage Efficiency Policy on the workload, the policies must have been created before you start creating the workload.

Steps

- On the **Provision Workload** page, add the name of the workload that you want to create, and then select the cluster from the available list on which you want to create the workload.
- Based on the cluster that you have selected, the **STORAGE VM** field filters the available storage virtual machines for that cluster. Select the required storage VM from the list.
Based on the SMB, NFS, iSCSI, and FCP services supported on the storage VM, the NAS and SAN options are enabled in the Host Information section.
- In the **Storage and Optimization** section, assign the storage capacity and Performance Service Level, and optionally, Storage Efficiency Policy for the workload.

The specifications for the Performance Service Level are applied to the workload, and the specifications for the Storage Efficiency Policy are assigned to the volume and LUN, as applicable, when the workload is created.

- Select the **Enforce performance limits** check box if you want to enforce the Performance Service Level that you have assigned to the workload. Assigning a Performance Service Level to a workload ensures that the aggregate on which the workload is created can support the performance and capacity objectives defined in the respective Performance Service Level. For example, if a workload is assigned "Extreme Performance", the aggregate on which the workload is to be provisioned should have the capability of supporting the performance and capacity objectives of the "Extreme Performance" Performance Service Level, such as SSD storage.



Unless you select this check box, the Performance Service Level is not applied on the workload, and the status of the workload on the dashboard appears as unassigned.

- If you are creating an SMB and/or NFS file share, select the **NAS** button. This button is enabled for selection only when the storage VM has the required services enabled.



If your SVM is enabled for both SMB and NFS services, you can select the **Share by NFS** and **Share by SMB** buttons and create a file share that supports both NFS and SMB protocols. If you want to create either an SMB or CIFS share, select only the respective button.

- For NFS file shares, specify the IP address of the host or network to access the file share volume. You can enter comma-separated values for multiple hosts.

On adding the host IP address, an internal check runs for matching the host details with the SVM and the export policy for that host is created, or in case there is an existing policy, it is reused. If there are several NFS shares created for the same host, then an available export policy for the same host with matching rules is reused for all the files shares. The function of specifying rules of individual policies or reusing policies by providing specific policy keys is available when you provision the NFS share by using APIs.

- b. For an SMB share, specify which users or user groups can access the SMB share and assign the required permissions. For each group of users, a new access control list (ACL) is generated during the file share creation.
6. For creating LUNs, select the **SAN** button. This button is enabled for selection only when the selected storage VM has the required services enabled.
 - a. Select the host OS.
 - b. Specify the host mapping for the LUN. You can assign existing initiator groups (igroups), or define and map new igroups to LUNs to control which initiators have access to LUNs.



If you create a new igroup while provisioning the LUN, you need to wait till the next discovery cycle (up to 15 minutes) for using it for the LUN creation. It is therefore recommended that you use an existing igroup from the list of available igroups.

If you want to create a new igroup, select the **Create a new initiator group** button, and then enter the required information to create the igroup.

7. Click **Save**.

The workload is added to the list of storage workloads.

Managing Performance Service Levels

A Performance Service Level enables you to define the performance and storage objectives for a workload. You can assign a Performance Service Level to a workload when initially creating the workload, or afterwards by editing the workload.

The management and monitoring of storage resources are based on Service Level Objectives (SLOs). SLOs are defined by service level agreements that are based on performance and capacity. In Unified Manager, SLOs refer to the Performance Service Level definitions of the applications that are running on NetApp storage. Storage services are differentiated based on the performance and utilization of the underlying resources. A Performance Service Level is a description of the storage service objectives. A Performance Service Level enables the storage provider to specify the performance and capacity objectives for the workload.

Unified Manager provides a few predefined (or canned) Performance Service Levels: Extreme Performance, Performance, and Value. The Extreme Performance, Performance, and Value Performance Service Levels are applicable for most of the common storage workloads in a data center. Unified Manager also offers three Performance Service Levels for database applications: Extreme for Database Logs, Extreme for Database Shared Data, and Extreme for Database Data. These are extremely high-performance Performance Service Levels that support bursty IOPS and are appropriate for database applications with the highest throughput demand. If these predefined Performance Service Levels do not meet your requirements, then you can create new Performance Service Levels based on the definition of the predefined Performance Service Levels.

You can access the Performance Service Levels from the **Policies > Performance Service Levels** page and by using the storage provider APIs. Managing storage workloads by assigning Performance Service Levels to

them is convenient as you do not have to individually manage the storage workloads. Any modifications can also be managed by reassigning another Performance Service Level rather than managing them individually.

You cannot modify a Performance Service Level that is system-defined or that is currently assigned to a workload. You cannot delete a PSL that is assigned to a workload, or if it is the only available Performance Service Level.

The Performance Service Levels page lists the available Performance Service Level policies and enables you to add, edit, and delete them. This page displays the following information:

Field	Description
Name	Name of the Performance Service Level.
Type	Whether the policy is system-defined or user-defined.
Expected IOPS	Minimum number of IOPS that an application is expected to perform on a LUN or file share. Expected IOPS specifies the minimum expected IOPS allocated, based on the storage object allocated size.
Peak IOPS	Maximum number of IOPS that an application can perform on a LUN or file share. Peak IOPS specifies the maximum possible IOPS allocated, based on the storage object allocated size or the storage object used size. Peak IOPS are based on an allocation policy. The allocation policy is either allocated-space or used-space. When the allocation policy is set to allocated-space, the peak IOPS is calculated based on the size of the storage object. When the allocation policy is set to used-space, the peak IOPS is calculated based on the amount of data stored in the storage object, taking into account storage efficiencies. By default, the allocation policy is set to used-space.

Field	Description
Absolute minimum IOPS	<p>The absolute minimum IOPS, which is used as an override when the expected IOPS is less than this value. The value can be between 400 and 1000. The range of absolute minimum IOPS is calculated as follows:</p> $\text{Minimum value} = 1000 / \text{expected latency}$ <p>The default values of the system-defined Performance Service Levels are computed as the following:</p> <ul style="list-style-type: none"> Extreme Performance: If expected IOPS $\geq 6144/\text{TB}$, then absolute minimum IOPS = 1000 Performance: If expected IOPS $\geq 2048/\text{TB}$ and $< 6144/\text{TB}$, then absolute minimum IOPS = 500 Value: If expected IOPS $\geq 128/\text{TB}$ and $< 2048/\text{TB}$, then absolute minimum IOPS = 75 <p>The default values of the system-defined database Performance Service Levels are computed as the following:</p> <ul style="list-style-type: none"> Extreme for Database Logs: If expected IOPS ≥ 22528, then absolute minimum IOPS = 4000 Extreme for Database Shared Data: If expected IOPS ≥ 16384, then absolute minimum IOPS = 2000 Extreme for Database Data: If expected IOPS ≥ 12288, then absolute minimum IOPS = 2000
Expected latency	Expected latency for storage IOPS in milliseconds per operation (ms/op).
Capacity	Total available and used capacity in the clusters.
Workloads	Number of storage workloads that have been assigned the Performance Service Level.

Note that if workloads exceed the expected latency value for 30% of the time during the previous hour that Unified Manager will generate one of the following events to notify you of a potential performance issue: “Workload Volume Latency Threshold Breached as defined by Performance Service Level Policy” or “Workload LUN Latency Threshold Breached as defined by Performance Service Level Policy”.

The following table provides information about the system-defined Performance Service Levels:

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Extreme Performance	Provides extremely high throughput at a very low latency Ideal for latency-sensitive applications	1	12288	6144	1000
Performance	Provides high throughput at a low latency Ideal for database and virtualized applications	2	4096	2048	500
Value	Provides high storage capacity and moderate latency Ideal for high-capacity applications such as email, web content, file shares, and backup targets	17	512	128	75

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Extreme for Database Logs	<p>Provides maximum throughput at the lowest latency.</p> <p>Ideal for database applications supporting database logs. This PSL provides the highest throughput because database logs are extremely bursty and logging is constantly in demand.</p>	1	45056	22528	4000
Extreme for Database Shared Data	<p>Provides very high throughput at the lowest latency.</p> <p>Ideal for database applications data that is stored in a common data store, but is shared across databases.</p>	1	32768	16384	2000
Extreme for Database Data	<p>Provides high throughput at the lowest latency.</p> <p>Ideal for database applications data, such as database table information and metadata.</p>	1	24576	12288	2000

Guidelines for creating a custom Performance Service Level

If the existing Performance Service Levels do not meet the Service Level Objective (SLO) requirements for your storage workloads, you can create a custom Performance Service Level. However, it is recommended that you attempt to use the system-defined Performance Service Levels for your storage workloads, and only create custom Performance Service Levels if necessary.

Creating and editing Performance Service Levels

When the system-defined Performance Service Levels do not match your workload requirements, you can create your own Performance Service Levels that are optimized for your workloads.

Before you begin

- You must have the Application Administrator role.
- The Performance Service Level name must be unique, and you cannot use the following reserved keywords:

Prime, Extreme, Performance, Value, Unassigned, Learning, Idle, Default, and None.

About this task

You create and edit custom Performance Service Levels from the Performance Service Levels page by defining the service level objectives you require for the applications that will access storage.



You cannot modify a Performance Service Level if it is currently assigned to a workload.

Steps

1. In the left navigation pane under **Settings**, select **Policies > Performance Service Levels**.
2. In the **Performance Service Levels** page, click the appropriate button depending on whether you want to create a new Performance Service Level or if you want to edit an existing Performance Service Level.

To...	Follow these steps...
Create a new Performance Service Level	Click Add .
Edit an existing Performance Service Level	Select an existing Performance Service Level, and then click Edit .

The page to add or edit a Performance Service Level is displayed.

3. Customize the Performance Service Level by specifying the performance objectives, and then click **Submit** to save the Performance Service Level.

After you finish

You can apply the new or changed Performance Service Level to workloads (LUNs, NFS File Shares, CIFS Shares) from the Workloads page or when provisioning a new workload.

Managing Storage Efficiency Policies

A Storage Efficiency Policy (SEP) enables you to define the storage efficiency characteristics of a workload. You can assign a Storage Efficiency Policy to a workload when initially creating the workload, or afterwards by editing the workload.

Storage efficiency includes using technologies, such as thin provisioning, deduplication, and data compression that increase storage utilization and decrease storage costs. While creating Storage Efficiency Policies, you can use these space-saving technologies either individually or together to achieve maximum storage efficiency. When you associate the policies with your storage workloads, the specified policy settings are assigned to them. Unified Manager enables you to assign system-defined and user-defined Storage Efficiency Policies to optimize storage resources in your data center.

Unified Manager provides two system-defined Storage Efficiency Policies: High and Low. These SEPs are applicable to most of the storage workloads in a data center, however, you can create your own policies if the system-defined SEPs do not meet your requirements.

You cannot modify a Storage Efficiency Policy that is system-defined or that is currently assigned to a workload. You cannot delete a Storage Efficiency Policy that is assigned to a workload, or if it is the only available Storage Efficiency Policy.

The Storage Efficiency Policies page lists the available Storage Efficiency Policies and enables you to add, edit, and delete customize SEPs. This page displays the following information:

Field	Description
Name	Name of the Storage Efficiency Policy.
Type	Whether the policy is system-defined or user-defined.
Space Reserve	Whether the volume is thin-provisioned or thick-provisioned.
Deduplication	Whether deduplication is enabled on the volume: <ul style="list-style-type: none">• Inline: Deduplication occurs while being written on the volume• Background: Deduplication occurs in the background• Disable: Deduplication is disabled on the volume

Field	Description
Compression	<p>Whether data compression is enabled on the volume:</p> <ul style="list-style-type: none"> • Inline: Data compression occurs while being written on the volume • Background: Data compression occurs in the background • Disable: Data compression is disabled on the volume
Workloads	Number of storage workloads that have been assigned the Storage Efficiency Policy

Guidelines for creating a custom Storage Efficiency Policy

If the existing Storage Efficiency Policies do not meet policy requirements for your storage workloads, you can create a custom Storage Efficiency Policy. However, it is recommended that you attempt to use the system-defined Storage Efficiency Policies for your storage workloads, and only create custom Storage Efficiency Policies if necessary.

Creating and editing Storage Efficiency Policies

When the system-defined Storage Efficiency Policies do not match your workload requirements, you can create your own Storage Efficiency Policies that are optimized for your workloads.

Before you begin

- You must have the Application Administrator role.
- The Storage Efficiency Policy name must be unique, and you cannot use the following reserved keywords: High, Low, Unassigned, Learning, Idle, Default, and None.

About this task

You create and edit custom Storage Efficiency Policies from the Storage Efficiency Policies page by defining the storage efficiency characteristics you require for the applications that will access storage.



You cannot modify a Storage Efficiency Policy if it is currently assigned to a workload.

Steps

1. In the left navigation pane under **Settings**, select **Policies > Storage Efficiency Policies**.
2. In the **Storage Efficiency Policies** page, click the appropriate button depending on whether you want to create a new Storage Efficiency Policy or if you want to edit an existing Storage Efficiency Policy.

To...	Follow these steps...
Create a new Storage Efficiency Policy	Click Add
Edit an existing Storage Efficiency Policy	Select an existing Storage Efficiency Policy and click Edit

The page to add or edit a Storage Efficiency Policy is displayed.

3. Customize the Storage Efficiency Policy by specifying the storage efficiency characteristics, and then click **Submit** to save the Storage Efficiency Policy.

After you finish

You can apply the new or changed Storage Efficiency Policy to workloads (LUNs, NFS File Shares, CIFS Shares) from the Workloads page or when provisioning a new workload.

Copyright information

Copyright © 2025 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.