



System: Storage array settings

SANtricity 11.6

NetApp
June 11, 2024

Table of Contents

- System: Storage array settings 1
 - Concepts 1
 - How tos 4
 - FAQs 12

System: Storage array settings

Concepts

Cache settings and performance

Cache memory is an area of temporary volatile storage on the controller that has a faster access time than the drive media.

With caching, overall I/O performance can be increased as follows:

- Data requested from the host for a read might already be in the cache from a previous operation, thus eliminating the need for drive access.
- Write data is written initially to the cache, which frees the application to continue instead of waiting for the data to be written to the drive.

The default cache settings meet the requirements for most environments, but you can change them if you want.

Storage array cache settings

For all volumes in the storage array, you can specify the following values from the System page:

- **Start value for flushing** — the percentage of unwritten data in the cache that triggers a cache flush (write to disk). When the cache holds the specified start percentage of unwritten data, a flush is triggered. By default, the controller starts flushing the cache when the cache reaches 80 percent full.
- **Cache block size** — the maximum size of each cache block, which is an organizational unit for cache management. The cache block size is by default 8 KiB, but can be set to 4, 8, 16, or 32 KiB. Ideally the cache block size should be set to the predominant I/O size of your applications. File systems or database applications generally use smaller sizes, while a larger size is good for applications requiring large data transfer or sequential I/O.

Volume cache settings

For individual volumes in a storage array, you can specify the following values from the Volumes page ([Storage](#) > [Volumes](#)):

- **Read caching** — The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.
 - **Dynamic read cache prefetch** — Dynamic cache read prefetch allows the controller to copy additional sequential data blocks into the cache while it is reading data blocks from a drive to the cache. This caching increases the chance that future requests for data can be filled from the cache. Dynamic cache read prefetch is important for multimedia applications that use sequential I/O. The rate and amount of data that is prefetched into cache is self-adjusting based on the rate and request size of the host reads. Random access does not cause data to be prefetched into cache. This feature does not apply when read caching is disabled.
- **Write caching** — The write cache is a buffer that stores data from the host that has not yet been written to the drives. The data stays in the write cache until it is written to the drives. Write caching can increase I/O performance.



Possible loss of data — If you enable the Write caching without batteries option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the Write caching without batteries option.

- **Write caching without batteries** — The write caching without batteries setting lets write caching continue even when the batteries are missing, failed, discharged completely, or not fully charged. Choosing write caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.
- **Write caching with mirroring** — Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.

Automatic load balancing overview

Automatic load balancing provides improved I/O resource management by reacting dynamically to load changes over time and automatically adjusting volume controller ownership to correct any load imbalance issues when workloads shift across the controllers.

The workload of each controller is continually monitored and, with cooperation from the multipath drivers installed on the hosts, can be automatically brought into balance whenever necessary. When workload is automatically re-balanced across the controllers, the storage administrator is relieved of the burden of manually adjusting volume controller ownership to accommodate load changes on the storage array.

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.

Enabling and disabling Automatic Load Balancing

Automatic Load Balancing is enabled by default on all storage arrays.

You might want to disable Automatic Load Balancing on your storage array for the following reasons:

- You do not want to automatically change a particular volume's controller ownership to balance workload.
- You are operating in a highly tuned environment where load distribution is purposefully set up to achieve a specific distribution between the controllers.

Host types that support the Automatic Load Balancing feature

Even though Automatic Load Balancing is enabled at the storage array level, the host type you select for a host or host cluster has a direct influence on how the feature operates.

When balancing the storage array's workload across controllers, the Automatic Load Balancing feature attempts to move volumes that are accessible by both controllers and that are mapped only to a host or host cluster capable of supporting the Automatic Load Balancing feature.

This behavior prevents a host from losing access to a volume due to the load balancing process; however, the presence of volumes mapped to hosts that do not support Automatic Load Balancing affects the storage array's ability to balance workload. For Automatic Load Balancing to balance the workload, the multipath driver must support TPGS and the host type must be included in the following table.



For a host cluster to be considered capable of Automatic Load Balancing, all hosts in that group must be capable of supporting Automatic Load Balancing.

Host type supporting Automatic Load Balancing	With this multipath driver
Windows or Windows Clustered	MPIO with NetApp E-Series DSM
Linux DM-MP (Kernel 3.10 or later)	DM-MP with <code>scsi_dh_alua</code> device handler
VMware	Native Multipathing Plugin (NMP) with <code>VMW_SATP_ALUA</code> Storage Array Type plug-in



With minor exceptions, host types that do not support Automatic Load Balancing continue to operate normally whether or not the feature is enabled. One exception is that if a system has a failover, storage arrays move unmapped or unassigned volumes back to the owning controller when the data path returns. Any volumes that are mapped or assigned to non-Automatic Load Balancing hosts are not moved.

See the [Interoperability Matrix Tool](#) for compatibility information for specific multipath driver, OS level, and controller-drive tray support.

Verifying OS compatibility with the Automatic Load Balancing feature

Verify OS compatibility with the Automatic Load Balancing feature before setting up a new (or migrating an existing) system.

1. Go to the [Interoperability Matrix Tool](#) to find your solution and verify support.

If your system is running Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11, contact technical support.

2. Update and configure the `/etc/multipath.conf` file.
3. Ensure that both `retain_attached_device_handler` and `detect_prio` are set to `yes` for the applicable vendor and product, or use default settings.

Default host operating system type

The default host type is used by the storage array when hosts are initially connected. It defines how the controllers in the storage array work with the host's operating system when volumes are accessed. You can change the host type if there is a need to change how the storage array operates, relative to the hosts that are connected to it.

Generally, you will change the default host type before you connect hosts to the storage array or when you connect additional hosts.

Keep these guidelines in mind:

- If all of the hosts you plan to connect to the storage array have the same operating system (homogenous host environment), then change the host type to match the operating system.
- If there are hosts with different operating systems that you plan to connect to the storage array (heterogeneous host environment), change the host type to match the majority of the hosts' operating systems.

For example, if you are connecting eight different hosts to the storage array, and six of those hosts are running a Windows operating system, you must select Windows as the default host operating system type.

- If the majority of the connected hosts have a mix of different operating systems, change the host type to Factory Default.

For example, if you are connecting eight different hosts to the storage array, and two of those hosts are running a Windows operating system, three are running a VMware operating system, and another three are running a Linux operating system, you must select Factory Default as the default host operating system type.

How tos

Edit storage array name

You can change the storage array name that appears in the title bar of SANtricity System Manager.

Steps

1. Select **Settings** > **System**.
2. Under **General**, look for the **Name:** field.

If a storage array name has not been defined, this field displays "Unknown."

3. Click the **Edit** (pencil) icon next to the storage array name.

The field becomes editable.

4. Enter a new name.

A name can contain letters, numbers, and the special characters underscore (), dash (-), and hash sign (#). A name cannot contain spaces. A name can have a maximum length of 30 characters. The name must be unique.

5. Click the **Save** (check mark) icon.



If you want to close the editable field without making changes, click the **Cancel** (X) icon.

Results

The new name appears in the title bar of SANtricity System Manager.

Turn on storage array locator lights

To find the physical location of a storage array in a cabinet, you can turn on its locator (LED) lights.

Steps

1. Select **Settings > System**.
2. Under **General**, click **Turn on Storage Array Locator Lights**.

The **Turn On Storage Array Locator Lights** dialog box opens, and the corresponding storage array's locator lights turn on.

3. When you have physically located the storage array, return to the dialog box and select **Turn Off**.

Results

The locator lights turn off, and the dialog box closes.

Synchronize storage array clocks

If Network Time Protocol (NTP) is not enabled, you can manually set the clocks on the controllers so they are synchronized with the management client (the system used to run the browser that accesses SANtricity System Manager).

About this task

Synchronization ensures that event time stamps in the event log match time stamps written to the host log files. During the synchronization process, the controllers remain available and operational.



If NTP is enabled in System Manager, do not use this option to synchronize clocks. Instead, NTP automatically synchronizes the clocks with an external host using SNTP (Simple Network Time Protocol).



After synchronization, you might notice that performance statistics are lost or skewed, schedules are impacted (ASUP, snapshots, etc.), and time stamps in log data are skewed. Using NTP avoids this problem.

Steps

1. Select **Settings > System**.
2. Under **General**, click **Synchronize Storage Array Clocks**.

The Synchronize Storage Array Clocks dialog box opens. It shows the current date and time for the controller(s) and the computer used as the management client.



For simplex storage arrays, only one controller is shown.

3. If the times shown in the dialog box do not match, click **Synchronize**.

Results

After synchronization is successful, event time stamps are the same for the event log and host logs.

Save storage array configuration

You can save a storage array's configuration information in a script file to save time setting up additional storage arrays with the same configuration.

Before you begin

The storage array must not be undergoing any operation that changes its logical configuration settings. Examples of these operations include creating or deleting volumes, downloading controller firmware, assigning or modifying hot spare drives, or adding capacity (drives) to a volume group.

About this task

Saving the storage array configuration generates a command line interface (CLI) script that contains storage array settings, volume configuration, host configuration, or host-to-volume assignments for a storage array. You can use this generated CLI script to replicate a configuration to another storage array with the exact same hardware configuration.

However, you should not use this generated CLI script for disaster recovery. Instead, to do a system restore, use the configuration database backup file that you create manually or contact technical support to get this data from the latest Auto-Support data.

This operation *does not* save these settings:

- The life of the battery
- The controller time-of-day
- The nonvolatile static random access memory (NVSRAM) settings
- Any premium features
- The storage array password
- The operating status and states of the hardware components
- The operating status (except Optimal) and states of the volume groups
- Copy services, such as mirroring and volume copy



Risk of application errors — Do not use this option if the storage array is undergoing an operation that will change any logical configuration setting. Examples of these operations include creating or deleting volumes, downloading controller firmware, assigning or modifying hot spare drives, or adding capacity (drives) to a volume group.

Steps

1. Select **Settings > System**.
2. Select **Save Storage Array Configuration**.
3. Select the items of the configuration that you want to save:
 - **Storage array settings**
 - **Volume configuration**
 - **Host configuration**
 - **Host-to-volume assignments**



If you select the **Host-to-volume assignments** item, the **Volume configuration** item and the **Host configuration** item are also selected by default. You cannot save **Host-to-volume assignments** without also saving **Volume configuration** and **Host configuration**.

4. Click **Save**.

The file is saved in the Downloads folder for your browser with the name `storage-array-configuration.cfg`.

After you finish

To load the saved storage array configuration onto another storage array, use the SANtricity command line interface (SMcli) with the `-f` option to apply the `.cfg` file.



You can also load a storage array configuration to other storage arrays by using the Unified Manager interface (select **Manage** > **Import Settings**).

Clear storage array configuration

Use the Clear Configuration operation when you want to delete all the pools, volume groups, volumes, host definitions, and host assignments from the storage array.

Before you begin

- Before clearing the storage array configuration, back up the data.

About this task

There are two Clear Storage Array Configuration options:

- **Volume** — Typically, you might use the Volume option to reconfigure a test storage array as a production storage array. For example, you might configure a storage array for testing, and then, when you are done testing, remove the test configuration and set up the storage array for a production environment.
- **Storage Array** — Typically, you might use the Storage Array option to move a storage array to another department or group. For example, you might be using a storage array in Engineering, and now Engineering is getting a new storage array, so you want to move the current storage array to Administration where it will be reconfigured.

The Storage Array option deletes some additional settings.

	Volume	Storage Array
Deletes pools and volume groups	X	X
Deletes volumes	X	X
Deletes hosts and host clusters	X	X
Deletes host assignments	X	X
Deletes storage array name		X

	Volume	Storage Array
Resets storage array cache settings to default		X



Risk of data loss — This operation deletes all data from your storage array. (It does not do a secure erase.) You cannot cancel this operation after it starts. Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Settings > System**.
2. Select **Clear Storage Array Configuration**.
3. In the drop-down list, select either **Volume** or **Storage Array**.
4. **Optional:** If you want to save the configuration (not the data), use the links in the dialog box.
5. Confirm that you want to perform the operation.

Results

- The current configuration is deleted, destroying all existing data on the storage array.
- All drives are unassigned.

Configure login banner

You can create a login banner that is presented to users before they establish sessions in SANtricity System Manager. The banner can include an advisory notice and a consent message.

About this task

When you create a banner, it appears before the login screen in a dialog box.

Steps

1. Select **Settings > System**.
2. Under the **General** section, select **Configure Login Banner**.

The Configure Login Banner dialog box opens.

3. Enter the text you want to appear in the login banner.



Do not use HTML or other markup tags for formatting.

4. Click **Save**.

Results

The next time users log in to System Manager, the text opens in a dialog box. Users must click **OK** to continue to the login screen.

Manage session timeouts

You can configure timeouts in SANtricity System Manager, so that users' inactive

sessions are disconnected after a specified time.

About this task

By default, the session timeout for System Manager is 30 minutes. You can adjust that time or you can disable session timeouts altogether.



If Access Management is configured using the Security Assertion Markup Language (SAML) capabilities embedded in the array, a session timeout might occur when the user's SSO session reaches its maximum limit. This might occur before the System Manager session timeout.

Steps

1. Select **Settings > System**.
2. Under the **General** section, select **Enable/Disable Session Timeout**.

The **Enable/Disable Session Timeout** dialog box opens.

3. Use the spinner controls to increase or decrease the time in minutes.

The minimum timeout you can set for System Manager is 15 minutes.



To disable session timeouts, deselect the **Set the length of time...** checkbox.

4. Click **Save**.

Change cache settings for the storage array

For all volumes in the storage array, you can adjust the cache memory settings for flushing and block size.

About this task

Cache memory is an area of temporary volatile storage on the controller, which has a faster access time than the drive media. To tune cache performance, you can adjust the following settings:

Cache setting	Description
Start demand cache flushing	Start demand cache flushing specifies the percentage of unwritten data in the cache that triggers a cache flush (write to disk). By default, cache flushing starts when unwritten data reaches 80% capacity. A higher percentage is a good choice for environments with primarily write operations, so new write requests can be processed by cache without having to go to the disk. Lower settings are better in environments where the I/O is erratic (with data bursts), so that the system flushes cache frequently between data bursts. However, a start percentage lower than 80% may cause decreased performance.

Cache setting	Description
Cache block size	The cache block size determines the maximum size of each cache block, which is an organizational unit for cache management. By default, the block size is 32 KiB. System Manager allows the cache block size to be 4, 8, 16, or 32 KiBs. Applications use different block sizes, which have an impact on storage performance. A smaller size is a good choice for file systems or database applications. A larger size is ideal for applications that generate sequential I/O, such as multimedia.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Cache Settings**.

The Change Cache Settings dialog box opens.

3. Adjust the following values:
 - **Start demand cache flushing** — Choose a percentage that is appropriate for the I/O used in your environment. If you choose a value lower than 80%, you may notice decreased performance.
 - **Cache block size** — Choose a size that is appropriate for your applications.
4. Click **Save**.

Set host connectivity reporting

You can enable host connectivity reporting so the storage array continuously monitors the connection between the controllers and the configured hosts, and then alerts you if the connection is disrupted. This feature is enabled by default.

About this task

If you disable host connectivity reporting, the system no longer monitors connectivity or multipath driver issues with a host connected to the storage array.



Disabling host connectivity reporting also disables automatic load balancing, which monitors and balances controller resource utilization.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Enable/Disable Host Connectivity Reporting**.

The text below this option indicates whether it is currently enabled or disabled.

A confirmation dialog box opens.

3. Click **Yes** to continue.

By selecting this option, you toggle the feature between enabled/disabled.

Set automatic load balancing

The Automatic Load Balancing feature ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers. This feature is enabled by default, but you can disable it from System Manager.

About this task

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.

You might want to disable Automatic Load Balancing on your storage array for the following reasons:

- You do not want to automatically change a particular volume's controller ownership to balance workload.
- You are operating in a highly tuned environment where load distribution is purposefully set up to achieve a specific distribution between the controllers.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Enable/Disable Automatic Load Balancing**.

The text below this option indicates whether the feature is currently enabled or disabled.

A confirmation dialog box opens.

3. Confirm by clicking **Yes** to continue.

By selecting this option, you toggle the feature between enabled/disabled.



If this feature is moved from disabled to enabled, the Host Connectivity Reporting feature is automatically enabled as well.

Change default host type

Use the Change Default Host Operating System setting to change the default host type at the storage array level. Generally, you will change the default host type before you connect hosts to the storage array or when you connect additional hosts.

About this task

Keep these guidelines in mind:

- If all of the hosts you plan to connect to the storage array have the same operating system (homogenous host environment), then change the host type to match the operating system.
- If there are hosts with different operating systems that you plan to connect to the storage array (heterogeneous host environment), change the host type to match the majority of the hosts' operating systems.

For example, if you are connecting eight different hosts to the storage array, and six of those hosts are

running a Windows operating system, you must select Windows as the default host operating system type.

- If the majority of the connected hosts have a mix of different operating systems, change the host type to Factory Default.

For example, if you are connecting eight different hosts to the storage array, and two of those hosts are running a Windows operating system, three are running a VMware operating system, and another three are running a Linux operating system, you must select Factory Default as the default host operating system type.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Default Host Operating System Type**.
3. Select the host operating system type that you want to use as the default.
4. Click **Change**.

Enable or disable legacy management interface

You can enable or disable the legacy management interface (SYMBOL), which is a method of communication between the storage array and the management client.

About this task

By default, the legacy management interface is on. If you disable it, the storage array and management client will use a more secure method of communication (REST API over https); however, certain tools and tasks might be affected if it is disabled.



For the EF600 storage system, this feature is disabled by default.

The setting affects operations as follows:

- **On** (default) — Required setting for configuring mirroring with the CLI and some other tools, such as the OCI adapter.
- **Off** — Required setting to enforce confidentiality in communications between the storage array and the management client, and to access external tools. Recommended setting when configuring a Directory Server (LDAP).

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Management Interface**.
3. In the dialog box, click **Yes** to continue.

FAQs

What is controller cache?

The controller cache is a physical memory space that streamlines two types of I/O (input/output) operations: between the controllers and hosts, and between the controllers and disks.

For read and write data transfers, the hosts and controllers communicate over high-speed connections. However, communications from the back-end of the controller to the disks is slower, because disks are relatively slow devices.

When the controller cache receives data, the controller acknowledges to the host applications that it is now holding the data. This way, the host applications do not need to wait for the I/O to be written to disk. Instead, applications can continue operations. The cached data is also readily accessible by server applications, eliminating the need for extra disk reads to access the data.

The controller cache affects the overall performance of the storage array in several ways:

- The cache acts as a buffer, so that host and disk data transfers do not need to be synchronized.
- The data for a read or write operation from the host might be in cache from a previous operation, which eliminates the need to access the disk.
- If write caching is used, the host can send subsequent write commands before the data from a previous write operation is written to disk.
- If cache prefetch is enabled, sequential read access is optimized. Cache prefetch makes a read operation more likely to find its data in the cache, instead of reading the data from disk.



Possible loss of data — If you enable the **Write caching without batteries** option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the **Write caching without batteries** option.

What is cache flushing?

When the amount of unwritten data in the cache reaches a certain level, the controller periodically writes cached data to a drive. This write process is called "flushing."

The controller uses two algorithms for flushing cache: demand-based and age-based. The controller uses a demand-based algorithm until the amount of cached data drops below the cache flush threshold. By default, a flush begins when 80 percent of the cache is in use.

In System Manager, you can set the "Start demand cache flushing" threshold to best support the type of I/O used in your environment. In an environment that is primarily write operations, you should set the "Start demand cache flushing" percentage high to increase the probability that any new write requests can be processed by cache without having to go to the disk. A high percentage setting limits the number of cache flushes so that more data remains in cache, which increases the chance of more cache hits.

In an environment where the I/O is erratic (with data bursts), you can use low cache flushing so that the system flushes cache frequently between data bursts. In a diverse I/O environment that processes a variety of loads, or when the type of loads are unknown, set the threshold at 50 percent as a good middle ground. Be aware that if you choose a start percentage lower than 80 percent, you might see decreased performance because data needed for a host read might not be available. Choosing a lower percentage also increases the number of disk writes necessary to maintain the cache level, which increases system overhead.

The age-based algorithm specifies the period of time during which write data can remain in the cache before it is eligible to be flushed to the disks. The controllers use the age-based algorithm until the cache flush threshold is reached. The default is 10 seconds, but this time period is counted only during periods of inactivity. You cannot modify the flush timing in System Manager; instead, you must use the **Set Storage Array** command in the command-line interface (CLI).



Possible loss of data — If you enable the **Write caching without batteries** option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the **Write caching without batteries** option.

What is cache block size?

The storage array's controller organizes its cache into "blocks," which are chunks of memory that can be 8, 16, 32 KiB in size. All volumes on the storage system share the same cache space; therefore, the volumes can have only one cache block size.

Applications use different block sizes, which can have an impact on storage performance. By default, the block size in System Manager is 32 KiB, but you can set the value to 8, 16, 32 KiBs. A smaller size is a good choice for file systems or database applications. A larger size is a good choice for applications that require large data transfer, sequential I/O, or high bandwidth, such as multimedia.

When should I synchronize storage array clocks?

You should manually synchronize the controller clocks in the storage array if you notice that the time stamps shown in System Manager are not aligned with time stamps shown in your management client (the computer that is accessing System Manager through the browser). This task is only necessary if NTP (Network Time Protocol) is not enabled in System Manager.



We highly recommend that you use an NTP server instead of manually synchronizing the clocks. NTP automatically synchronizes the clocks with an external server using SNTP (Simple Network Time Protocol).

You can check synchronization status from the Synchronize Storage Array Clocks dialog box, which is available from the System page. If the times shown in the dialog box do not match, run a synchronization. You can periodically view this dialog box, which indicates whether the controller clocks' time displays have drifted apart and are no longer synchronized.

What is host connectivity reporting?

When host connectivity reporting is enabled, the storage array continuously monitors the connection between the controllers and the configured hosts, and then alerts you if the connection is disrupted.

Disruptions to the connection might occur if there is a loose, damaged, or missing cable, or another problem with the host. In these situations, the system might open a Recovery Guru message:

- **Host Redundancy Lost** — Opens if either controller cannot communicate with the host.
- **Host Type Incorrect** — Opens if the host's type is incorrectly specified on the storage array, which could result in failover problems.

You might want to disable host connectivity reporting in situations where rebooting a controller might take longer than the connection timeout. Disabling this feature suppresses Recovery Gurus messages.



Disabling host connectivity reporting also disables automatic load balancing, which monitors and balances controller resource use. However, if you re-enable host connectivity reporting, the automatic load balancing feature is not automatically re-enabled.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.