



# **Configure audit message and log destinations**

StorageGRID 11.8

NetApp  
March 19, 2024

# Table of Contents

- Configure audit message and log destinations ..... 1
- Considerations for using an external syslog server ..... 1
- Configure audit messages and external syslog server ..... 6

# Configure audit message and log destinations

## Considerations for using an external syslog server

An external syslog server is a server outside of StorageGRID you can use to collect system audit information in a single location. Using an external syslog server enables you to reduce network traffic on your Admin Nodes and manage the information more efficiently. For StorageGRID, the outbound syslog message packet format is compliant with RFC 3164.

The types of audit information you can send to the external syslog server include:

- Audit logs containing the audit messages generated during normal system operation
- Security-related events such as logins and escalations to root
- Application logs that might be requested if it is necessary to open a support case to troubleshoot an issue you have encountered

## When to use an external syslog server

An external syslog server is especially useful if you have a large grid, use multiple types of S3 applications, or want to retain all audit data. Sending audit information to an external syslog server enables you to:

- Collect and manage audit information such as audit messages, application logs, and security events more efficiently.
- Reduce network traffic on your Admin Nodes because audit information is transferred directly from the various Storage Nodes to the external syslog server, without having to go through an Admin Node.



When logs are sent to an external syslog server, single logs greater than 8,192 bytes are truncated at the end of the message to conform with common limitations in external syslog server implementations.



To maximize the options for full data recovery in the event of a failure of the external syslog server, up to 20 GB of local logs of audit records (`localaudit.log`) are maintained on each node.

## How to configure an external syslog server

To learn how to configure an external syslog server, see [Configure audit messages and external syslog server](#).

If you plan to configure use the TLS or RELP/TLS protocol, you must have the following certificates:

- **Server CA certificates:** One or more trusted CA certificates for verifying the external syslog server in PEM encoding. If omitted, the default Grid CA certificate will be used.
- **Client certificate:** The client certificate for authentication to the external syslog server in PEM encoding.
- **Client private key:** Private key for the client certificate in PEM encoding.



If you use a client certificate you must also use a client private key. If you provide an encrypted private key, you must also provide the passphrase. There is no significant security benefit from using an encrypted private key because the key and passphrase must be stored; using an unencrypted private key, if available, is recommended for simplicity.

## How to estimate the size of the external syslog server

Normally, your grid is sized to achieve a required throughput, defined in terms of S3 operations per second or bytes per second. For example, you might have a requirement that your grid handle 1,000 S3 operations per second, or 2,000 MB per second, of object ingests and retrievals. You should size your external syslog server according to your grid's data requirements.

This section provides some heuristic formulas that help you estimate the rate and average size of log messages of various types that your external syslog server needs to be capable of handling, expressed in terms of the known or desired performance characteristics of the grid (S3 operations per second).

### Use S3 operations per second in estimation formulas

If your grid was sized for a throughput expressed in bytes per second, you must convert this sizing into S3 operations per second to use the estimation formulas. To convert grid throughput, you must first determine your average object size, which you can do using the information in existing audit logs and metrics (if any), or by using your knowledge of the applications that will use StorageGRID. For example, if your grid was sized to achieve a throughput of 2,000 MB/second, and your average object size is 2 MB, then your grid was sized to be able to handle 1,000 S3 operations per second (2,000 MB / 2 MB).



The formulas for external syslog server sizing in the following sections provide common-case estimates (rather than worst-case estimates). Depending on your configuration and workload, you might see a higher or lower rate of syslog messages or volume of syslog data than the formulas predict. The formulas are meant to be used as guidelines only.

### Estimation formulas for audit logs

If you have no information about your S3 workload other than number of S3 operations per second your grid is expected to support, you can estimate the volume of audit logs your external syslog server will need to handle using the following formulas, under the assumption that you leave the Audit Levels set to the default values (all categories set to Normal, except Storage, which is set to Error):

```
Audit Log Rate = 2 x S3 Operations Rate  
Audit Log Average Size = 800 bytes
```

For example, if your grid is sized for 1,000 S3 operations per second, your external syslog server should be sized to support 2,000 syslog messages per second and should be able to receive (and typically store) audit log data at a rate of 1.6 MB per second.

If you know more about your workload, more accurate estimations are possible. For audit logs, the most important additional variables are the percentage of S3 operations that are PUTs (vs. GETS), and the average size, in bytes, of the following S3 fields (4-character abbreviations used in the table are audit log field names):

Code	Field	Description
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
S3BK	S3 bucket	The S3 bucket name.
S3KY	S3 key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.

Let's use  $P$  to represent the percentage of S3 operations that are PUTs, where  $0 \leq P \leq 1$  (so for a 100% PUT workload,  $P = 1$ , and for a 100% GET workload,  $P = 0$ ).

Let's use  $K$  to represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like /my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then the value of  $K$  is 90 (13+13+28+36).

If you can determine values for  $P$  and  $K$ , you can estimate the volume of audit logs your external syslog server will need to handle using the following formulas, under the assumption that you leave the Audit Levels set to the defaults (all categories set to Normal, except Storage, which is set to Error):

$$\text{Audit Log Rate} = ((2 \times P) + (1 - P)) \times \text{S3 Operations Rate}$$

$$\text{Audit Log Average Size} = (570 + K) \text{ bytes}$$

For example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 1,500 syslog messages per second and should be able to receive (and typically store) audit log data at a rate of approximately 1 MB per second.

### Estimation formulas for non-default audit levels

The formulas provided for audit logs assume the use of default audit level settings (all categories set to Normal, except Storage, which is set to Error). Detailed formulas for estimating the rate and average size of audit messages for non-default audit level settings aren't available. However, the following table can be used to make a rough estimate of the rate; you can use the average size formula provided for audit logs, but be aware that it is likely to result in an over-estimate because the "extra" audit messages are, on average, smaller than the default audit messages.

Condition	Formula
Replication: Audit levels all set to Debug or Normal	Audit log rate = 8 x S3 Operations Rate

Condition	Formula
Erasure coding: audit levels all set to Debug or Normal	Use same formula as for default settings

### Estimation formulas for security events

Security events aren't correlated with S3 operations and typically produce a negligible volume of logs and data. For these reasons, no estimation formulas are provided.

### Estimation formulas for application logs

If you have no information about your S3 workload other than the number of S3 operations per second your grid is expected to support, you can estimate the volume of applications logs your external syslog server will need to handle using the following formulas:

```
Application Log Rate = 3.3 x S3 Operations Rate
Application Log Average Size = 350 bytes
```

So, for example, if your grid is sized for 1,000 S3 operations per second, your external syslog server should be sized to support 3,300 application logs per second and be able to receive (and store) application log data at a rate of about 1.2 MB per second.

If you know more about your workload, more accurate estimations are possible. For application logs, the most important additional variables are the data protection strategy (replication vs. erasure coding), the percentage of S3 operations that are PUTs (vs. GETs/other), and the average size, in bytes, of the following S3 fields (4-character abbreviations used in table are audit log field names):

Code	Field	Description
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
S3BK	S3 bucket	The S3 bucket name.
S3KY	S3 key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.

### Example sizing estimations

This section explains example cases of how to use the estimation formulas for grids with the following methods of data protection:

- Replication
- Erasure coding

### If you use replication for data protection

Let P represent the percentage of S3 operations that are PUTs, where  $0 \leq P \leq 1$  (so for a 100% PUT workload,  $P = 1$ , and for a 100% GET workload,  $P = 0$ ).

Let K represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like /my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then K has a value of 90 (13+13+28+36).

If you can determine values for P and K, you can estimate the volume of application logs your external syslog server will have to be able to handle using the following formulas.

```
Application Log Rate = ((1.1 x P) + (2.5 x (1 - P))) x S3 Operations Rate
Application Log Average Size = (P x (220 + K)) + ((1 - P) x (240 + (0.2 x K))) Bytes
```

So, for example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 1800 application logs per second, and will be receiving (and typically storing) application data at a rate of 0.5 MB per second.

### If you use erasure coding for data protection

Let P represent the percentage of S3 operations that are PUTs, where  $0 \leq P \leq 1$  (so for a 100% PUT workload,  $P = 1$ , and for a 100% GET workload,  $P = 0$ ).

Let K represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like /my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then K has a value of 90 (13+13+28+36).

If you can determine values for P and K, you can estimate the volume of application logs your external syslog server will have to be able to handle using the following formulas.

```
Application Log Rate = ((3.2 x P) + (1.3 x (1 - P))) x S3 Operations Rate
Application Log Average Size = (P x (240 + (0.4 x K))) + ((1 - P) x (185 + (0.9 x K))) Bytes
```

So, for example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 2,250 application logs per second and should be able to receive (and typically store) application data at a rate of 0.6 MB per second.

# Configure audit messages and external syslog server

You can configure a number of settings related to audit messages. You can adjust the number of audit messages recorded; define any HTTP request headers you want to include in client read and write audit messages; configure an external syslog server; and specify where audit logs, security event logs, and StorageGRID software logs are sent.

Audit messages and logs record system activities and security events, and are essential tools for monitoring and troubleshooting. All StorageGRID nodes generate audit messages and logs to track system activity and events.

Optionally, you can configure an external syslog server to save audit information remotely. Using an external server minimizes the performance impact of audit message logging without reducing the completeness of audit data. An external syslog server is especially useful if you have a large grid, use multiple types of S3 applications, or want to retain all audit data. See [Considerations for external syslog server](#) for details.

## Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Maintenance or Root access permission](#).
- If you plan to configure an external syslog server, you have reviewed the [considerations for using an external syslog server](#) and ensured that the server has enough capacity to receive and store the log files.
- If you plan to configure an external syslog server using TLS or RELP/TLS protocol, you have the required server CA and client certificates and the client private key.

## Change audit message levels

You can set a different audit level for each of the following categories of messages in the audit log:

Audit category	Default setting	More information
System	Normal	<a href="#">System audit messages</a>
Storage	Error	<a href="#">Object storage audit messages</a>
Management	Normal	<a href="#">Management audit message</a>
Client reads	Normal	<a href="#">Client read audit messages</a>
Client writes	Normal	<a href="#">Client write audit messages</a>
ILM	Normal	<a href="#">ILM audit messages</a>
Cross-grid replication	Error	<a href="#">CGRR: Cross-Grid Replication Request</a>



These defaults apply if you initially installed StorageGRID using version 10.3 or later. If you initially used an earlier version of StorageGRID, the default for all categories is set to Normal.





During upgrades, audit level configurations will not be effective immediately.

### Steps

1. Select **CONFIGURATION > Monitoring > Audit and syslog server**.
2. For each category of audit message, select an audit level from the drop-down list:

Audit level	Description
Off	No audit messages from the category are logged.
Error	Only error messages are logged—audit messages for which the result code was not "successful" (SUCS).
Normal	Standard transactional messages are logged—the messages listed in these instructions for the category.
Debug	Deprecated. This level behaves the same as the Normal audit level.

The messages included for any particular level include those that would be logged at the higher levels. For example, the Normal level includes all of the Error messages.



If you don't require a detailed record of client read operations for your S3 applications, optionally change the **Client Reads** setting to **Error** to decrease the number of audit messages recorded in the audit log.

3. Select **Save**.

A green banner indicates your configuration has been saved.

## Define HTTP request headers

You can optionally define any HTTP request headers you want to include in client read and write audit messages. These protocol headers apply to S3 and Swift requests only.

### Steps

1. In the **Audit protocol headers** section, define the HTTP request headers you want to include in client read and write audit messages.

Use an asterisk (\*) as a wildcard to match zero or more characters. Use the escape sequence (\\*) to match a literal asterisk.

2. Select **Add another header** to create additional headers, if needed.

When HTTP headers are found in a request, they are included in the audit message under the field HTRH.



Audit protocol request headers are logged only if the audit level for **Client Reads** or **Client Writes** is not **Off**.

3. Select **Save**

A green banner indicates your configuration has been saved.

## Use an external syslog server

You can optionally configure an external syslog server to save audit logs, application logs, and security event logs to a location outside of your grid.



If you don't want to use an external syslog server, skip this step and go to [Select audit information destinations](#).



If the configuration options available in this procedure aren't flexible enough to meet your requirements, additional configuration options can be applied using the `audit-destinations` endpoints, which are in the private API section of the [Grid Management API](#). For example, you can use the API if you want to use different syslog servers for different groups of nodes.

### Enter syslog information

Access the Configure external syslog server wizard and provide the information StorageGRID needs to access the external syslog server.

#### Steps

1. From the Audit and syslog server page, select **Configure external syslog server**. Or, if you have previously configured an external syslog server, select **Edit external syslog server**.

The Configure external syslog server wizard appears.

2. For the **Enter syslog info** step of the wizard, enter a valid fully qualified domain name or an IPv4 or IPv6 address for the external syslog server in the **Host** field.
3. Enter the destination port on the external syslog server (must be an integer between 1 and 65535). The default port is 514.
4. Select the protocol used to send audit information to the external syslog server.

Using **TLS** or **RELPL/TLS** is recommended. You must upload a server certificate to use either of these options. Using certificates helps secure the connections between your grid and the external syslog server. For more information, see [Manage security certificates](#).

All protocol options require support by, and configuration of, the external syslog server. You must choose an option that is compatible with the external syslog server.



Reliable Event Logging Protocol (RELPL) extends the functionality of the syslog protocol to provide reliable delivery of event messages. Using RELPL can help prevent the loss of audit information if your external syslog server has to restart.

5. Select **Continue**.
6. If you selected **TLS** or **RELPL/TLS**, upload the server CA certificates, client certificate, and client private key.
  - a. Select **Browse** for the certificate or key you want to use.
  - b. Select the certificate or key file.
  - c. Select **Open** to upload the file.

A green check appears next to the certificate or key file name, notifying you that it has been uploaded successfully.

7. Select **Continue**.

## Manage syslog content

You can select which information to send to the external syslog server.

### Steps

1. For the **Manage syslog content** step of the wizard, select each type of audit information you want to send to the external syslog server.

- **Send audit logs:** Sends StorageGRID events and system activities
- **Send security events:** Sends security events such as when an unauthorized user attempts to sign in or a user signs in as root
- **Send application logs:** Sends log files useful for troubleshooting including:
  - `bycast-err.log`
  - `bycast.log`
  - `jaeger.log`
  - `nms.log` (Admin Nodes only)
  - `prometheus.log`
  - `raft.log`
  - `hagroups.log`

For information about StorageGRID software logs, see [StorageGRID software logs](#).

2. Use the drop-down menus to select the severity and facility (type of message) for each category of audit information you want to send.

Setting severity and facility values can help you aggregate the logs in customizable ways for easier analysis.

a. For **Severity**, select **Passthrough**, or select a severity value between 0 and 7.

If you select a value, the selected value will be applied to all messages of this type. Information about different severities will be lost if you override severity with a fixed value.

Severity	Description
Passthrough	Each message sent to the external syslog to have the same severity value as when it was logged locally onto the node: <ul style="list-style-type: none"> <li>• For audit logs, the severity is "info."</li> <li>• For security events, the severity values are generated by the Linux distribution on the nodes.</li> <li>• For application logs, the severities vary between "info" and "notice," depending on what the issue is. For example, adding an NTP server and configuring an HA group gives a value of "info," while intentionally stopping the SSM or RSM service gives a value of "notice."</li> </ul>
0	Emergency: System is unusable
1	Alert: Action must be taken immediately
2	Critical: Critical conditions
3	Error: Error conditions
4	Warning: Warning conditions
5	Notice: Normal but significant condition
6	Informational: Informational messages
7	Debug: Debug-level messages

b. For **Facility**, select **Passthrough**, or select a facility value between 0 and 23.

If you select a value, it will be applied to all messages of this type. Information about different facilities will be lost if you override facility with a fixed value.

Facility	Description
Passthrough	<p>Each message sent to the external syslog to have the same facility value as when it was logged locally onto the node:</p> <ul style="list-style-type: none"> <li>• For audit logs, the facility sent to the external syslog server is "local7."</li> <li>• For security events, the facility values are generated by the linux distribution on the nodes.</li> <li>• For application logs, the application logs sent to the external syslog server have the following facility values: <ul style="list-style-type: none"> <li>◦ <code>bycast.log</code>: user or daemon</li> <li>◦ <code>bycast-err.log</code>: user, daemon, local3, or local4</li> <li>◦ <code>jaeger.log</code>: local2</li> <li>◦ <code>nms.log</code>: local3</li> <li>◦ <code>prometheus.log</code>: local4</li> <li>◦ <code>raft.log</code>: local5</li> <li>◦ <code>hagroups.log</code>: local6</li> </ul> </li> </ul>
0	kern (kernel messages)
1	user (user-level messages)
2	mail
3	daemon (system daemons)
4	auth (security/authorization messages)
5	syslog (messages generated internally by syslogd)
6	lpr (line printer subsystem)
7	news (network news subsystem)
8	UUCP
9	cron (clock daemon)
10	security (security/authorization messages)
11	FTP
12	NTP

Facility	Description
13	logaudit (log audit)
14	logalert (log alert)
15	clock (clock daemon)
16	local0
17	local1
18	local2
19	local3
20	local4
21	local5
22	local6
23	local7

3. Select **Continue**.

### Send test messages

Before starting to use an external syslog server, you should request that all nodes in your grid send test messages to the external syslog server. You should use these test messages to help you validate your entire log collection infrastructure before you commit to sending data to the external syslog server.



Don't use the external syslog server configuration until you confirm that the external syslog server received a test message from each node in your grid and that the message was processed as expected.

### Steps

1. If you don't want to send test messages because you are certain your external syslog server is configured properly and can receive audit information from all the nodes in your grid, select **Skip and finish**.

A green banner indicates that the configuration has been saved.

2. Otherwise, select **Send test messages** (recommended).

Test results continuously appear on the page until you stop the test. While the test is in progress, your audit messages continue to be sent to your previously configured destinations.

3. If you receive any errors, correct them and select **Send test messages** again.

See [Troubleshoot an external syslog server](#) to help you resolve any errors.

4. Wait until you see a green banner indicating all nodes have passed testing.
5. Check your syslog server to determine if test messages are being received and processed as expected.



If you are using UDP, check your entire log collection infrastructure. The UDP protocol does not allow for as rigorous error detection as the other protocols.

6. Select **Stop and finish**.

You are returned to the **Audit and syslog server** page. A green banner indicates that the syslog server configuration has been saved.



StorageGRID audit information is not sent to the external syslog server until you select a destination that includes the external syslog server.

## Select audit information destinations

You can specify where audit logs, security event logs, and [StorageGRID software logs](#) are sent.



Some destination are available only if you have configured an external syslog server.

### Steps

1. On the Audit and syslog server page, select the destination for audit information.



**Local nodes only** and **External syslog server** typically provide better performance.

Option	Description
Local nodes only	Audit messages, security event logs, and application logs are not sent to Admin Nodes. Instead, they are saved only on the nodes that generated them ("the local node"). The audit information generated on every local node is stored in <code>/var/local/log/localaudit.log</code>  <b>Note:</b> StorageGRID periodically removes local logs in a rotation to free up space. When the log file for a node reaches 1 GB, the existing file is saved, and a new log file is started. The rotation limit for the log is 21 files. When the 22nd version of the log file is created, the oldest log file is deleted. On average about 20 GB of log data is stored on each node.
Admin Nodes/local nodes	Audit messages are sent to the audit log ( <code>/var/local/log/audit.log</code> ) on Admin Nodes, and security event logs and application logs are stored on the nodes that generated them.

Option	Description
External syslog server	Audit information is sent to an external syslog server and saved on the local nodes. The type of information sent depends upon how you configured the external syslog server. This option is enabled only after you have configured an external syslog server.
Admin Node and external syslog server	Audit messages are sent to the audit log (/var/local/log/audit.log) on Admin Nodes, and audit information is sent to the external syslog server and saved on the local node. The type of information sent depends upon how you configured the external syslog server. This option is enabled only after you have configured an external syslog server.

2. Select **Save**.

A warning message appears.

3. Select **OK** to confirm that you want to change the destination for audit information.

A green banner indicates that the audit configuration has been saved.

New logs are sent to the destinations you selected. Existing logs remain in their current location.



## 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.