



TR-4626: Load balancers

StorageGRID solutions and resources

NetApp
November 21, 2025

Table of Contents

TR-4626: Load balancers	1
Use third-party load balancers with StorageGRID	1
Load balancer basics	1
Where to find additional information	2
Use StorageGRID load balancers	2
StorageGRID Gateway Node load balancer versus third-party load balancer	2
Learn how to implement SSL certificates for HTTPS in StorageGRID	3
Configure trusted third-party load balancer in StorageGRID	4
Learn about local traffic manager load balancers	4
Create a resource group of Storage Nodes	4
Health check	5
Virtual hosted-style addressing	6
SSL termination	6
Pass through configuration	7
Source client IP visibility	7
Learn about few use cases for StorageGRID configurations	7
F5 BIG-IP local traffic manager health check monitor for S3 bucket	7
Citrix NetScaler	9
Loadbalancer.org	9
Kemp	9
HAProxy	9
Validate SSL connection in StorageGRID	10
Understand global load balancing requirements for StorageGRID	10

TR-4626: Load balancers

Use third-party load balancers with StorageGRID

Learn about the role of a third-party and global load balancers in an object storage systems like StorageGRID.

General guidance for implementing NetApp® StorageGRID® with third-party load balancers.


Object storage is synonymous with the term cloud storage, and, as you would expect, applications that leverage cloud storage address that storage through a URL. Behind that simple URL, StorageGRID can scale capacity, performance, and durability in a single site or over geo-distributed sites. The component that makes this simplicity possible is a load balancer.

The purpose of this document is to educate StorageGRID customers about load balancer options and provide general guidance for the configuration of third-party load balancers.

Load balancer basics

Load balancers are an essential component of an enterprise grade object storage system such as StorageGRID. StorageGRID consists of multiple storage nodes, each of which can present the entire Simple Storage Service (S3) name space for a given StorageGRID instance. Load balancers create a highly available endpoint behind which we can place StorageGRID nodes. StorageGRID is unique among S3-compatible object storage systems in that it provides its own load balancer, but it also supports third-party or general-purpose load balancers such as F5, Citrix Netscaler, HA Proxy, NGINX, and so on.

The following figure uses the example URL/ fully qualified domain name (FQDN) “s3.company.com”. The load balancer creates a virtual IP (VIP) that resolves to the FQDN through DNS, then directs any requests from applications to a pool of StorageGRID nodes. The load balancer performs a health check on each node and only establishes connections to healthy nodes.



The figure shows the StorageGRID provided load balancer, but the concept is the same for third-party load balancers. Applications establish an HTTP session using the VIP on the load balancer and the traffic passes through the load balancer to the storage nodes. By default, all traffic, from application to load balancer, and from load balancer to storage node is encrypted through HTTPS. HTTP is a supported option.

Local and global load balancers

There are two types of load balancers:

- **Local Traffic Managers (LTM).** Spreads connections over a pool of nodes in a single site.
- **Global Service Load Balancer (GSLB).** Spreads connections over multiple sites, effectively load balancing LTM load balancers. Think of a GSLB as an intelligent DNS server. When a client requests a StorageGRID endpoint URL, the GSLB resolves it to the VIP of an LTM based on availability or other factors (for example, which site can provide lower latency to the application). While an LTM is always required, a GSLB is optional depending on the number of StorageGRID sites and your application requirements.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp StorageGRID Documentation Center
<https://docs.netapp.com/us-en/storagegrid/>
- NetApp StorageGRID Enablement
<https://docs.netapp.com/us-en/storagegrid-enable/>
- StorageGRID f5 load balancer design considerations
<https://www.netapp.com/blog/storagegrid-f5-load-balancer-design-considerations/>
- Loadbalancer.org—Load balancing NetApp StorageGRID
<https://www.loadbalancer.org/applications/load-balancing-netapp-storagegrid/>
- Kemp—Load balancing NetApp StorageGRID
<https://support.kemptechnologies.com/hc/en-us/articles/360045186451-NetApp-StorageGRID>

Use StorageGRID load balancers

Learn about the role of a StorageGRID Gateway Node load balancer.

General guidance for implementing NetApp® StorageGRID® Gateway Nodes.

StorageGRID Gateway Node load balancer versus third-party load balancer

StorageGRID is unique among S3-compatible object storage vendors in that it provides a native load balancer available as a purpose-built appliance, VM, or container. The StorageGRID provided load balancer is also referred to as a Gateway Node.

For customers that do not already own a load balancer such as F5, Citrix, and so on, implementation of a third-party load balancer can be very complex. The StorageGRID load balancer greatly simplifies load balancer operations.

The Gateway Node is an enterprise grade, highly available, and high-performance load balancer. Customers can choose to implement the Gateway Node, third-party load balancer, or even both, in the same grid. The Gateway Node is a local traffic manager versus a GSLB.

The StorageGRID load balancer provides the following advantages:

- **Simplicity.** Automatic configuration of resource pools, health checks, patching, and maintenance, all managed by StorageGRID.
- **Performance.** The StorageGRID load balancer is dedicated to StorageGRID, can provide high performance caching, and you do not compete with other applications for bandwidth.

- **Cost.** The virtual machine (VM) and container versions are provided at no additional cost.
- **Traffic classifications.** The Advanced Traffic Classification feature allows for StorageGRID-specific QoS rules along with workload analytics.
- **Future StorageGRID specific features.** StorageGRID will continue to optimize and add innovative features to the load balancer over upcoming releases.

As an integrated node of StorageGRID, the local traffic manager has the ability to use advanced health checking to distribute requests based on Storage Node health, load and resource availability. In addition it has the ability to distribute the load across multiple sites when the StorageGRID link costs are set to "0" between the sites. In the event the Storage Nodes are unavailable but the Gateway Node is available in a site, the load will automatically be directed to another site in the grid.

The Load balancer caching feature of the Gateway Node is intended to provide a substantial performance improvement for certain workloads (such as AI training) which re-read a data set multiple times as part of processing that data.

Caching gateway nodes can also be deployed physically distant from the rest of the grid enabling better performance and lower WAN network utilization in some workloads. The cache operates in a read back mode where writes are not cached and do not modify the state of the cache. Each Caching Gateway Node operates independently of any other caching Gateway node.

For details about deploying the StorageGRID Gateway Node, see the [StorageGRID documentation](#).

Learn how to implement SSL certificates for HTTPS in StorageGRID

Understand the importance and the steps to implement of SSL certificates in StorageGRID.

If you are using HTTPs, you must have a Secure Sockets Layer (SSL) certificate. The SSL protocol identifies the clients and endpoints, validating them as trusted. SSL also provides encryption of the traffic. The SSL certificate must be trusted by the clients. To accomplish this, the SSL certificate can be from a globally trusted Certificate Authority (CA) such as DigiCert, a private CA running in your infrastructure, or a self-signed certificate generated by the host.

Using a globally trusted CA certificate is the preferred method as there is no additional client-side actions required. The certificate is loaded into the load balancer or StorageGRID, and the clients trust and connect to the endpoint.

Using a private CA requires the root and all subordinate certificates be added to the client. The process to trust a private CA certificate can vary by client operating system and applications. For example, in ONTAP for FabricPool, you must upload each certificate in the chain individually (root certificate, subordinate certificate, endpoint certificate) to the ONTAP cluster.

Using a self-signed certificate requires the client to trust the provided certificate without any CA to verify the authenticity. Some applications might not accept self-signed certificates and have no ability to ignore verification.

The placement of the SSL certificate in the client load balancer StorageGRID path depends on where you need the SSL termination to be. You can configure a load balancer to be the termination endpoint for the client, and then re-encrypt or hot encrypt with a new SSL certificate for the load balancer to StorageGRID connection. Or you can pass through the traffic and let StorageGRID be the SSL termination endpoint. If the load balancer is the SSL termination endpoint, the certificate is installed on the load balancer and contains the subject name

for the DNS name/URL and any alternative URL/DNS names for which a client is configured to connect to the StorageGRID target through the load balancer, including any wild card names. If the load balancer is configured for pass through, the SSL certificate must be installed in StorageGRID. Again, the certificate must contain the subject name for the DNS name/URL, and any alternative URL/DNS names for which a client is configured to connect to the StorageGRID target through the load balancer, including any wild card names. Individual Storage Node names do not need to be included on the certificate, only the endpoint URLs.

```
Subject DN: /C=US/postalCode=94089/ST=California/L=Sunnyvale/street=495 East Java Dr/O=NetApp, Inc./OU=IT1/OU=Unified Communication  
s/CN=webscaledemo.netapp.com  
Serial Number: 37:4C:6B:51:61:84:50:F8:7A:29:D9:83:24:12:36:2C  
Issuer DN: /C=GB/ST=Greater Manchester/L=Salford/O=Sectigo Limited/CN=Sectigo RSA Organization Validation Secure Server CA  
Issued On: 2019-05-23T00:00:00.000Z  
Expires On: 2021-05-22T23:59:59.000Z  
Alternative Names: DNS:webscaledemo.netapp.com  
DNS:*.webscaledemo-rtp.netapp.com  
DNS:*.webscaledemo.netapp.com  
DNS:webscaledemo-rtp.netapp.com  
SHA-1 Fingerprint: 60:91:44:E5:4F:7E:25:6B:B5:A0:19:87:D1:F2:8C:DD:AD:3A:88:CD  
SHA-256 Fingerprint: FE:21:5D:BF:08:D9:5A:E5:09:CF:F6:3F:D3:5C:1E:9B:33:63:63:CA:25:2D:3F:39:0B:6A:B8:EC:08:BC:57:43
```

Configure trusted third-party load balancer in StorageGRID

Learn how to configure trusted third-party load balancer in StorageGRID.

If you are using one or more external layer 7 load balancers, and an S3 bucket or group policies that are IP based, StorageGRID must determine the real sender's IP address. It does this by looking at the X-Forwarded-For (XFF) header, which is inserted into the request by the load balancer. As the XFF header can be easily spoofed in requests sent directly to the Storage Nodes, StorageGRID must confirm that each request is being routed by a trusted layer 7 load balancer. If StorageGRID cannot trust the source of the request, it will ignore the XFF header. There is a Grid Management API to allow a list of trusted external layer 7 load balancers to be configured. This new API is private and is subject to change in future StorageGRID releases. For the most up to date information, see the KB article, [How to configure StorageGRID to work with third-party Layer 7 load balancers](#).

Learn about local traffic manager load balancers

Explore the guidance for local traffic manager load balancers and determine the optimal configuration.

The following is presented as general guidance for configuration of third-party load balancers. Work with your load balancer administrator to determine the optimal configuration for your environment.

Create a resource group of Storage Nodes

Group StorageGRID Storage Nodes into a resource pool or service group (the terminology might differ with specific load balancers).

StorageGRID Storage Nodes present the S3 API on the following ports:

- S3 HTTPS: 18082
- S3 HTTP: 18084

Most customers choose to present the APIs on the virtual server through the standard HTTPS and HTTP ports (443 and 80).



Each StorageGRID site requires a default of three Storage Nodes, two of which must be healthy.

Health check

Third-party load balancers require a method to determine the health of each node and its eligibility to receive traffic. NetApp recommends the HTTP OPTIONS method to perform the health check. The load balancer issues HTTP OPTIONS requests to each individual Storage Node and expects a 200 status response.

If any Storage Node does not provide a 200 response, that node is not able to service storage requests. Your application and business requirements should determine the timeout for these checks and the action your load balancer takes.

For example, if three of four Storage Nodes in data center 1 are down, you might direct all traffic to data center 2.

The recommended polling interval is once per second, marking the node offline after three failed checks.

S3 health check example

In the following example, we send OPTIONS and check for 200 OK. We use OPTIONS because Amazon S3) does not support unauthorized requests.

```
curl -X OPTIONS https://10.63.174.75:18082 --verbose --insecure
* Rebuilt URL to: https://10.63.174.75:18082/
*   Trying 10.63.174.75...
* TCP_NODELAY set
* Connected to 10.63.174.75 (10.63.174.75) port 18082 (#0)
* TLS 1.2 connection using TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
* Server certificate: webscale.st1.netapp.com
* Server certificate: NetApp Corp Issuing CA 1
* Server certificate: NetApp Corp Root CA
> OPTIONS / HTTP/1.1
> Host: 10.63.174.75:18082
> User-Agent: curl/7.51.0
> Accept: /
>
< HTTP/1.1 200 OK
< Date: Mon, 22 May 2017 15:17:30 GMT
< Connection: KEEP-ALIVE
< Server: StorageGRID/10.4.0
< x-amz-request-id: 3023514741
```

File or content-based health checks

In general, NetApp does not recommend file-based health checks. Typically, a small file — `healthcheck.htm`, for example — is created in a bucket with a read-only policy. This file is then fetched and evaluated by the load balancer. This approach has several disadvantages:

- **Dependent on a single account.** If the account that owns the file is disabled, the health check fails, and no storage requests are processed.
- **Data protection rules.** The default data protection scheme is a two-copy approach. In this scenario, if the two storage nodes hosting the health check file are unavailable, the health check fails, and storage requests are not sent to healthy storage nodes, rendering the grid offline.
- **Audit log bloat.** The load balancer fetches the file from every storage node every X minutes, creating many audit log entries.
- **Resource intensive.** Fetching the health check file from every node every few seconds consumes grid and network resources.

If a content-based health check is required, use a dedicated tenant with a dedicated S3 bucket.

Session persistence

Session persistence, or stickiness, refers to the time a given HTTP session is allowed to persist. By default, sessions are dropped by Storage Nodes after 10 minutes. Longer persistence can lead to better performance because applications do not have to reestablish their sessions for every action; however, holding these sessions open consumes resources. If you determine that your workload will benefit, you can reduce the session persistence on a third-party load balancer.

Virtual hosted-style addressing

Virtual hosted-style is now the default method for AWS S3, and while StorageGRID and many applications still support path style, it is best practice to implement virtual hosted-style support. Virtual hosted-style requests have the bucket as part of the host name.

To support virtual hosted-style, do the following:

- Support wildcard DNS lookups: *.s3.company.com
- Use an SSL certificate with subject alt names to support wildcard: *.s3.company.com
Some customers have expressed security concerns around the use of wildcard certificates. StorageGRID continues to support path style access, as do key applications such as FabricPool. That said, certain S3 API calls fail or behave improperly without virtual hosted support.

SSL termination

There are security benefits to SSL termination on third-party load balancers. If the load balancer is compromised, the grid is compartmentalized.

There are three supported configurations:

- **SSL pass-through.** The SSL certificate is installed on StorageGRID as a custom server certificate.
- **SSL termination and re-encryption (recommended).** This might be beneficial if you are already doing SSL certificate management on the load balancer rather than installing the SSL certificate on StorageGRID. This configuration provides the additional security benefit of limiting the attack surface to the load balancer.
- **SSL termination with HTTP.** In this configuration, SSL is terminated on the third-party load balancer and communication from the load balancer to StorageGRID is nonencrypted to take advantage of SSL off-load (with SSL libraries embedded in modern processors this is of limited benefit).

Pass through configuration

If you prefer to configure your load balancer for pass through, you must install the certificate on StorageGRID. Go to **Configuration > Server Certificates > Object Storage API Service Endpoints Server Certificate**.

Source client IP visibility

StorageGRID 11.4 introduced the concept of a trusted third-party load balancer. In order to forward the client application IP to StorageGRID, you must configure this feature. For more information, see [How to configure StorageGRID to work with third-party Layer 7 load balancers](#).

To enable the XFF header to be used to view the IP of the client application, follow these steps:

Steps

1. Record the client IP in the audit log.
2. Use `aws:SourceIp` S3 bucket or group policy.

Load balancing strategies

Most load balancing solutions offer multiple strategies for load balancing. The following are common strategies:

- **Round robin.** A universal fit but suffers with few nodes and large transfers clogging single nodes.
- **Least connection.** A good fit for small and mixed object workloads, resulting in an equal distribution of the connections to all nodes.

The choice of algorithm becomes less important with an increasing number of Storage Nodes to choose from.

Data path

All data flows through local traffic manager load balancers. StorageGRID does not support direct server routing (DSR).

Verifying distribution of connections

To verify that your method is distributing the load evenly across Storage Nodes, check the established sessions on each node in a given site:

- **UI Method.** Go to **Support > Metrics > S3 Overview > LDR HTTP Sessions**
- **Metrics API.** Use `storagegrid_http_sessions_incoming_currently_established`

Learn about few use cases for StorageGRID configurations

Explore few use cases for StorageGRID configurations implemented by customers and NetApp IT.

The following examples illustrate configurations as implemented by StorageGRID customers, including NetApp IT.

F5 BIG-IP local traffic manager health check monitor for S3 bucket

To configure the F5 BIG-IP local traffic manager health check monitor, follow these steps:

Steps

1. Create a new monitor.
 - a. In the Type field, enter HTTPS.
 - b. Configure the interval and timeout as desired.
 - c. In the Send String field, enter OPTIONS / HTTP/1.1\r\n\r\n.\r\n\r\n are carriage returns; different versions of BIG-IP software require zero, one, or two sets of \r\n sequences. For more information, see <https://support.f5.com/csp/article/K10655>.
 - d. In the Receive String field, enter: HTTP/1.1 200 OK.

General Properties	
Name	https_storagegrid
Description	
Type	HTTPS
Parent Monitor	https

Configuration: Basic	
Interval	5 seconds
Timeout	16 seconds
Send String	OPTIONS / HTTP/1.1\r\n\r\n
Receive String	HTTP/1.1 200 OK
Receive Disable String	
Cipher List	DEFAULT+SHA+3DES+KEDH
User Name	
Password	
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No
Alias Address	* All Addresses
Alias Service Port	* All Ports
Adaptive	<input type="checkbox"/> Enabled

2. In Create Pool, create one pool for each port required.
 - a. Assign the health monitor that you created in the previous step.
 - b. Select a load-balancing method.

c. Select service port: 18082 (S3).


d. Add nodes.

Citrix NetScaler

Citrix NetScaler creates a virtual server for the storage endpoint and refers to StorageGRID Storage Nodes as Application Servers, which are then grouped into Services.

Use the HTTPS-ECV health check monitor to create a custom monitor to perform the recommended health check by using the OPTIONS request and receiving 200. HTTP-ECV is configured with a send string and validates a receive string.

For more information, see the Citrix documentation, [Sample configuration for HTTP-ECV health check monitor](#).



The screenshot shows the 'Monitors' configuration screen in Citrix NetScaler. At the top, there are buttons for 'Add Binding', 'Edit Binding', 'Delete', and 'Edit Monitor'. A table lists a single monitor entry:

Monitor Name	Weight	State
STORAGE-GRID-TCP-ECV-MON	1	✓

Below the table, the 'Configure Monitor' section is expanded. It includes fields for 'Name' (set to 'STORAGE-GRID-TCP-ECV-MON'), 'Type' (set to 'TCP-ECV'), and 'Basic Parameters'. Under 'Basic Parameters', the 'Interval' is set to 5 seconds and the 'Response Timeout' is set to 2 seconds. In the 'Send String' section, the value 'OPTIONS / HTTP/1.1 r/r/r/r/n' is entered. In the 'Receive String' section, the value 'HTTP/2.0 200 OK' is entered. At the bottom of the configuration section, there are checkboxes for 'Secure' and 'SSL Profile', and buttons for 'Add' and 'Edit'.

Loadbalancer.org

Loadbalancer.org has conducted their own integration testing with StorageGRID and has an extensive configuration guide: https://pdfs.loadbalancer.org/NetApp_StorageGRID_Deployment_Guide.pdf.

Kemp

Kemp has conducted their own integration testing with StorageGRID and has an extensive configuration guide: <https://kemptechnologies.com/solutions/netapp/>.

HAProxy

Configure HAProxy to use the OPTIONS request and check for a 200 status response for the health check in haproxy.cfg. You can change the bind port in the front end to a different port, such as 443.

The following is an example for SSL termination on HAProxy:

```
frontend s3
    bind *:443 crt /etc/ssl/server.pem ssl
    default_backend s3-serve
rs
backend s3-servers
    balance leastconn
    option httpchk
    http-check expect status 200
    server dc1-s1 10.63.174.71:18082 ssl verify none check inter 3000
    server dc1-s2 10.63.174.72:18082 ssl verify none check inter 3000
    server dc1-s3 10.63.174.73:18082 ssl verify none check inter 3000
```

The following is an example for SSL pass-through:

```
frontend s3
    mode tcp
    bind *:443
    default_backend s3-servers
backend s3-servers
    balance leastconn
    option httpchk
    http-check expect status 200
    server dc1-s1 10.63.174.71:18082 check-ssl verify none inter 3000
    server dc1-s2 10.63.174.72:18082 check-ssl verify none inter 3000
    server dc1-s3 10.63.174.73:18082 check-ssl verify none inter 3000
```

For full examples of configurations for StorageGRID, see [Examples for HAProxy Configuration](#) on GitHub.

Validate SSL connection in StorageGRID

Learn how to validate the SSL connection in StorageGRID.


After your load balancer is configured, you should validate the connection using tools such as OpenSSL and the AWS CLI. Other applications, such as S3 Browser, might ignore SSL misconfiguration.

Understand global load balancing requirements for StorageGRID

Explore the design considerations and requirements for global load balancing in StorageGRID.

Global load balancing requires integrating with DNS to provide intelligent routing across multiple StorageGRID sites. This function falls outside of the StorageGRID domain and must be provided by a third-party solution

such as the load balancer products discussed previously and/or a DNS traffic control solution such as Infoblox. This top level load balancing provides smart routing to the closest destination site in the namespace, as well as outage detection and redirection to the next site in the namespace. A typical GSLB implementation consists of the top level GSLB with site pools containing site-local load balancers. The site load balancers contain pools of the local site Storage Nodes. This can include a combination of third-party load balancers for GSLB functions and StorageGRID providing the site-local load balancing, or a combination of third parties, or many of the third parties discussed previously can provide both GSLB and site-local load balancing.



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—with 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.