



Use FabricPool setup wizard

StorageGRID 11.8

NetApp
March 19, 2024

Table of Contents

- Use FabricPool setup wizard 1
 - Use FabricPool setup wizard: Considerations and requirements 1
 - Access and complete the FabricPool setup wizard 3

Use FabricPool setup wizard

Use FabricPool setup wizard: Considerations and requirements

You can use the FabricPool setup wizard to configure StorageGRID as the object storage system for a FabricPool cloud tier. After you complete the setup wizard, you can enter the required details into ONTAP System Manager.

When to use the FabricPool setup wizard

The FabricPool setup wizard guides you through each step of configuring StorageGRID for use with FabricPool and automatically configures certain entities for you, such as the ILM and traffic classification policies. As part of completing the wizard, you download a file that you can use to enter values into ONTAP System Manager. Use the wizard to configure your system more quickly and to make sure your settings conform to StorageGRID and FabricPool best practices.

Assuming you have Root access permission, you can complete the FabricPool setup wizard when you start using the StorageGRID Grid Manager, or you can access and complete the wizard at any later time. Depending on your requirements, you can also configure some or all of the required items manually and then use the wizard to assemble the values that ONTAP needs into a single file.



Use the FabricPool setup wizard unless you know you have special requirements or your implementation will require significant customization.

Before using the wizard

Confirm you have completed these prerequisite steps.

Review best practices

- You have a general understanding of the [information needed to attach StorageGRID as a cloud tier](#).
- You have reviewed the FabricPool best practices for:
 - [High availability \(HA\) groups](#)
 - [Load balancing](#)
 - [ILM rules and policy](#)

Obtain IP addresses and set up VLAN interfaces

If you will configure an HA group, you know which nodes ONTAP will connect to and which StorageGRID network will be used. You also know which values to enter for the subnet CIDR, gateway IP address, and virtual IP (VIP) addresses.

If you plan to use a virtual LAN to segregate FabricPool traffic, you have already configured the VLAN interface. See [Configure VLAN interfaces](#).

Configure identity federation and SSO

If you plan to use identity federation or single sign-on (SSO) for your StorageGRID system, you have enabled these features. You also know which federated group should have root access for the tenant account that ONTAP will use. See [Use identity federation](#) and [Configure single sign-on](#).

Obtain and configure domain names

- You know which fully qualified domain name (FQDN) to use for StorageGRID. Domain name server (DNS) entries will map this FQDN to the virtual IP (VIP) addresses of the HA group that you create using the wizard. See [Configure DNS server](#).
- If you plan to use S3 virtual hosted-style requests, you have [configured S3 endpoint domain names](#). ONTAP uses path-style URLs by default, but using virtual hosted-style requests is recommended.

Review load balancer and security certificate requirements

If you plan to use the StorageGRID load balancer, you have reviewed the general [considerations for load balancing](#). You have the certificates you will upload or the values you need to generate a certificate.

If you plan to use an external (third-party) load balancer endpoint, you have the fully qualified domain name (FQDN), port, and certificate for that load balancer.

Confirm ILM storage pool configuration

if you initially installed StorageGRID 11.6 or earlier, you have configured the storage pool you will use. In general, you should create a storage pool for each StorageGRID site you will use to store ONTAP data.



This prerequisite does not apply if you initially installed StorageGRID 11.7 or 11.8. When you initially install either of these versions, storage pools are automatically created for each site.

Relationship between ONTAP and the StorageGRID cloud tier

The FabricPool wizard guides you through the process of creating a single StorageGRID cloud tier that includes one StorageGRID tenant, one set of access keys, and one StorageGRID bucket. You can attach this StorageGRID cloud tier to one or more ONTAP local tiers.

Attaching a single cloud tier to multiple local tiers in a cluster is the general best practice. However, depending on your requirements, you might want to use more than one bucket or even more than one StorageGRID tenant for the local tiers in a single cluster. Using different buckets and tenants allows you to isolate data and data access between ONTAP local tiers, but is somewhat more complex to configure and manage.

NetApp does not recommend attaching a single cloud tier to local tiers in multiple clusters.



For the best practices for using StorageGRID with NetApp MetroCluster™ and FabricPool Mirror, see [TR-4598: FabricPool Best Practices in ONTAP](#).

Optional: Use a different bucket for each local tier

To use more than one bucket for the local tiers in an ONTAP cluster, add more than one StorageGRID cloud tier in ONTAP. Each cloud tier shares the same HA group, load balancer endpoint, tenant, and access keys, but uses a different container (StorageGRID bucket). Follow these general steps:

1. From StorageGRID Grid Manager, complete the FabricPool setup wizard for the first cloud tier.

2. From ONTAP System Manager, add a cloud tier and use the file you downloaded from StorageGRID to provide the required values.
3. From StorageGRID Tenant Manager, sign in to the tenant that was created by the wizard, and create a second bucket.
4. Complete the FabricPool wizard again. Select the existing HA group, load balancer endpoint, and tenant. Then, select the new bucket you created manually. Create a new ILM rule for the new bucket and activate an ILM policy to include that rule.
5. From ONTAP, add a second cloud tier but provide the new bucket name.

Optional: Use a different tenant and bucket for each local tier

To use more than one tenant and different sets of access keys for the local tiers in an ONTAP cluster, add more than one StorageGRID cloud tier in ONTAP. Each cloud tier shares the same HA group, load balancer endpoint, but uses a different tenant, access keys, and container (StorageGRID bucket). Follow these general steps:

1. From StorageGRID Grid Manager, complete the FabricPool setup wizard for the first cloud tier.
2. From ONTAP System Manager, add a cloud tier and use the file you downloaded from StorageGRID to provide the required values.
3. Complete the FabricPool wizard again. Select the existing HA group and load balancer endpoint. Create a new tenant and bucket. Create a new ILM rule for the new bucket and activate an ILM policy to include that rule.
4. From ONTAP, add a second cloud tier but provide the new access key, secret key, and bucket name.

Access and complete the FabricPool setup wizard

You can use the FabricPool setup wizard to configure StorageGRID as the object storage system for a FabricPool cloud tier.

Before you begin

- You have reviewed the [considerations and requirements](#) for using the FabricPool setup wizard.



If you want to configure StorageGRID for use with any other S3 client application, go to [Use S3 setup wizard](#).


- You have the [Root access permission](#).

Access the wizard

You can complete the FabricPool setup wizard when you start using the StorageGRID Grid Manager, or you can access and complete the wizard at any later time.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. If the **FabricPool and S3 setup wizard** banner appears on the dashboard, select the link in the banner. If the banner no longer appears, select the help icon from the header bar in the Grid Manager and select **FabricPool and S3 setup wizard**.



3. In the FabricPool section of the FabricPool and S3 setup wizard page, select **Configure now**.

Step 1 of 9: Configure HA group appears.

Step 1 of 9: Configure HA group

A high availability (HA) group is a collection of nodes that each contain the StorageGRID Load Balancer service. An HA group can contain Gateway Nodes, Admin Nodes, or both.

You can use an HA group to help keep FabricPool data connections available. An HA group uses virtual IP addresses (VIPs) to provide highly available access to the Load Balancer service. If the active interface in the HA group fails, a backup interface can manage the workload with little impact to FabricPool operations

For details about this task, see [Manage high availability groups](#) and [Best practices for high availability groups](#).

Steps

1. If you plan to use an external load balancer, you don't need to create an HA group. Select **Skip this step** and go to [Step 2 of 9: Configure load balancer endpoint](#).
2. To use the StorageGRID load balancer, create a new HA group or use an existing HA group.

Create HA group

- a. To create a new HA group, select **Create HA group**.
- b. For the **Enter details** step, complete the following fields.

Field	Description
HA group name	A unique display name for this HA group.
Description (optional)	The description of this HA group.

- c. For the **Add interfaces** step, select the node interfaces you want to use in this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

You can select one or more nodes, but you can select only one interface for each node.

- d. For the **Prioritize interfaces** step, determine the Primary interface and any backup interfaces for this HA group.

Drag rows to change the values in the **Priority order** column.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the active interface fails, the virtual IP (VIP) addresses move to the first backup interface in the priority order. If that interface fails, the VIP addresses move to the next backup interface, and so on. When failures are resolved, the VIP addresses move back to highest priority interface available.

- e. For the **Enter IP addresses** step, complete the following fields.

Field	Description
Subnet CIDR	The address of the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32). The network address must not have any host bits set. For example, 192.16.0.0/22.
Gateway IP address (optional)	Optional. If the ONTAP IP addresses used to access StorageGRID aren't on the same subnet as the StorageGRID VIP addresses, enter the StorageGRID VIP local gateway IP address. The local gateway IP address must be within the VIP subnet.
Virtual IP address	Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet and all will be active at the same time on the active interface. At least one address must be IPv4. Optionally, you can specify additional IPv4 and IPv6 addresses.

- f. Select **Create HA group** and then select **Finish** to return to the FabricPool setup wizard.
- g. Select **Continue** to go to the load balancer step.

Use existing HA group

- a. To use an existing HA group, select the HA group name from the **Select an HA group** drop-down list.
- b. Select **Continue** to go to the load balancer step.

Step 2 of 9: Configure load balancer endpoint

StorageGRID uses a load balancer to manage the workload from client applications, such as FabricPool. Load balancing maximizes speed and connection capacity across multiple Storage Nodes.

You can use the StorageGRID Load Balancer service, which exists on all Gateway and Admin Nodes, or you can connect to an external (third-party) load balancer. Using the StorageGRID load balancer is recommended.

For details about this task, see the general [considerations for load balancing](#) and the [best practices for load balancing for FabricPool](#).

Steps

1. Select or create a StorageGRID load balancer endpoint or use an external load balancer.

Create endpoint

- a. Select **Create endpoint**.
- b. For the **Enter endpoint details** step, complete the following fields.

Field	Description
Name	A descriptive name for the endpoint.
Port	<p>The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes, because these ports are reserved on Admin Nodes.</p> <p>Note: Ports used by other grid services aren't permitted. See the Network port reference.</p>
Client type	Must be S3 .
Network protocol	<p>Select HTTPS.</p> <p>Note: Communicating with StorageGRID without TLS encryption is supported but not recommended.</p>

-
-
- c. For the **Select binding mode** step, specify the binding mode. The binding mode controls how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Mode	Description
Global (default)	<p>Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN.</p> <p>Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.</p>
Virtual IPs of HA groups	<p>Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint.</p> <p>Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.</p>
Node interfaces	<p>Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.</p>
Node type	<p>Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.</p>

d. For the **Tenant access** step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets. Allow all tenants is almost always the appropriate option for the load balancer endpoint used for FabricPool. You must select this option if you are using the FabricPool setup wizard for a new StorageGRID system and you have not yet created any tenant accounts.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

e. For the **Attach certificate** step, select one of the following:

Field	Description
Upload certificate (recommended)	Use this option to upload a CA-signed server certificate, certificate private key, and optional CA bundle.
Generate certificate	Use this option to generate a self-signed certificate. See Configure load balancer endpoints for details of what to enter.
Use StorageGRID S3 and Swift certificate	This option is available only if you have already uploaded or generated a custom version of the StorageGRID global certificate. See Configure S3 and Swift API certificates for details.

f. Select **Finish** to return to the FabricPool setup wizard.

g. Select **Continue** to go to the tenant and bucket step.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

Use existing load balancer endpoint

- Select the name of an existing endpoint from the **Select a load balancer endpoint** drop-down list.
- Select **Continue** to go to the tenant and bucket step.

Use external load balancer

- Complete the following fields for the external load balancer.

Field	Description
FQDN	The fully qualified domain name (FQDN) of the external load balancer.
Port	The port number that FabricPool will use to connect to the external load balancer.
Certificate	Copy the server certificate for the external load balancer and paste it into this field.

- b. Select **Continue** to go to the tenant and bucket step.

Step 3 of 9: Tenant and bucket

A tenant is an entity that can use S3 applications to store and retrieve objects in StorageGRID. Each tenant has its own users, access keys, buckets, objects, and a specific set of capabilities. You must create a StorageGRID tenant before you can create the bucket that FabricPool will use.

A bucket is a container used to store a tenant's objects and object metadata. Although some tenants might have many buckets, the wizard lets you create or select only one tenant and one bucket at a time. You can use the Tenant Manager later to add any additional buckets you need.

You can create a new tenant and bucket for FabricPool use, or you can select an existing tenant and bucket. If you create a new tenant, the system automatically creates the access key ID and secret access key for the tenant's root user.

For details about this task, see [Create a tenant account for FabricPool](#) and [Create an S3 bucket and obtain an access key](#).

Steps

Create a new tenant and bucket or select an existing tenant.

New tenant and bucket

1. To create a new tenant and bucket, enter a **Tenant name**. For example, `FabricPool` tenant.
2. Define root access for the tenant account, based on whether your StorageGRID system uses [identity federation](#), [single sign-on \(SSO\)](#), or both.

Option	Do this
If identity federation is not enabled	Specify the password to use when signing into the tenant as the local root user.
If identity federation is enabled	<ol style="list-style-type: none">1. Select an existing federated group to have Root access permission for the tenant.2. Optionally, specify the password to use when signing in to the tenant as the local root user.
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

3. For **Bucket name**, enter the name of the bucket FabricPool will use to store ONTAP data. For example, `fabricpool-bucket`.



You can't change the bucket name after creating the bucket.

4. Select the **Region** for this bucket.

Use the default region (`us-east-1`) unless you expect to use ILM in the future to filter objects based on the bucket's region.

5. Select **Create and Continue** to create the tenant and bucket and to go to the download data step

Select tenant and bucket

The existing tenant account must have at least one bucket that does not have versioning enabled. You can't select an existing tenant account if no bucket exists for that tenant.

1. Select the existing tenant from the **Tenant name** drop-down list.
2. Select the existing bucket from the **Bucket name** drop-down list.

FabricPool does not support object versioning, so buckets that have versioning enabled aren't shown.




Don't select a bucket that has S3 Object Lock enabled for use with FabricPool.

3. Select **Continue** to go to the download data step.

Step 4 of 9: Download ONTAP settings

During this step, you download a file that you can use to enter values into ONTAP System Manager.

Steps

1. Optionally, select the copy icon () to copy both the access key ID and secret access key to the clipboard.

These values are included in the download file, but you might want to save them separately.

2. Select **Download ONTAP settings** to download a text file that contains the values you've entered so far.

The `ONTAP_FabricPool_settings_bucketname.txt` file includes the information you need to configure StorageGRID as the object storage system for a FabricPool cloud tier, including:

- Load balancer connection details, including the server name (FQDN), port, and certificate
- Bucket name
- Access key ID and secret access key for the root user of the tenant account

3. Save the copied keys and downloaded file to a secure location.



Don't close this page until you have copied both access keys, downloaded the ONTAP settings, or both. The keys will not be available after you close this page. Make sure to save this information in a secure location because it can be used to obtain data from your StorageGRID system.

4. Select the checkbox to confirm you have downloaded or copied the access key ID and secret access key.
5. Select **Continue** to go to the ILM storage pool step.

Step 5 of 9: Select a storage pool

A storage pool is a group of Storage Nodes. When you select a storage pool, you determine which nodes StorageGRID will use to store the data tiered from ONTAP.

For details about this step, see [Create a storage pool](#).

Steps

1. From the **Site** drop-down list, select the StorageGRID site you want to use for the data tiered from ONTAP.
2. From the **Storage pool** drop-down list, select the storage pool for that site.

The storage pool for a site includes all Storage Nodes at that site.

3. Select **Continue** to go to the ILM rule step.

Step 6 of 9: Review ILM rule for FabricPool

Information lifecycle management (ILM) rules control the placement, duration, and ingest behavior for all objects in your StorageGRID system.

The FabricPool setup wizard automatically creates the recommended ILM rule for FabricPool use. This rule applies only to the bucket you specified. It uses 2+1 erasure coding at a single site to store the data that is tiered from ONTAP.

For details about this step, see [Create ILM rule](#) and [Best practices for using ILM with FabricPool data](#).

Steps

1. Review the rule details.

Field	Description
Rule name	Automatically generated and can't be changed
Description	Automatically generated and can't be changed
Filter	The bucket name This rule only applies to objects that are saved in the bucket you specified.
Reference time	Ingest time The placement instruction starts when objects are initially saved to the bucket.
Placement instruction	Use 2+1 erasure coding

2. Sort the retention diagram by **Time period** and **Storage pool** to confirm the placement instruction.
 - The **Time period** for the rule is **Day 0 - forever**. **Day 0** means that the rule is applied when data is tiered from ONTAP. **Forever** means that StorageGRID ILM will not delete data that has been tiered from ONTAP.
 - The **Storage pool** for the rule is the storage pool you selected. **EC 2+1** means the data will be stored using 2+1 erasure coding. Each object will be saved as two data fragments and one parity fragment. The three fragments for each object will be saved to different Storage Nodes at a single site.
3. Select **Create and Continue** to create this rule and to go to the ILM policy step.

Step 7 of 9: Review and activate ILM policy

After the FabricPool setup wizard creates the ILM rule for FabricPool use, it creates an ILM policy. You must carefully simulate and review this policy before activating it.

For details about this step, see [Create ILM policy](#) and [Best practices for using ILM with FabricPool data](#).



When you activate a new ILM policy, StorageGRID uses that policy to manage the placement, duration, and data protection of all objects in the grid, including existing objects and newly ingested objects. In some cases, activating a new policy can cause existing objects to be moved to new locations.



To avoid data loss, do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period to **forever** to ensure that FabricPool objects aren't deleted by StorageGRID ILM.

Steps

1. Optionally, update the system-generated **Policy name**. By default, the system appends "+ FabricPool" to the name of your active or inactive policy, but you can provide your own name.

2. Review the list of rules in the inactive policy.

- If your grid doesn't have an inactive ILM policy, the wizard creates an inactive policy by cloning your active policy and adding the new rule to the top.
- If your grid already has an inactive ILM policy and that policy uses the same rules and same order as the active ILM policy, the wizard adds the new rule to the top of the inactive policy.
- If your inactive policy contains different rules or a different order than the active policy, the wizard creates a new inactive policy by cloning your active policy and adding the new rule to the top.

3. Review the order of the rules in the new inactive policy.

Because the FabricPool rule is the first rule, any objects in the FabricPool bucket are placed before the other rules in the policy are evaluated. Objects in any other buckets are placed by subsequent rules in the policy.

4. Review the retention diagram to learn how different objects will be retained.

- Select **Expand all** to see a retention diagram for each rule in the inactive policy.
- Select **Time period** and **Storage pool** to review the retention diagram. Confirm that any rules that apply to the FabricPool bucket or tenant retain objects **forever**.

5. When you have reviewed the inactive policy, select **Activate and continue** to activate the policy and go to the traffic classification step.



Errors in an ILM policy can cause irreparable data loss. Review the policy carefully before activating.

Step 8 of 9: Create traffic classification policy

As an option, the FabricPool setup wizard can create a traffic classification policy that you can use to monitor the FabricPool workload. The system-created policy uses a matching rule to identify all network traffic related to the bucket you created. This policy monitors traffic only; it does not limit traffic for FabricPool or any other clients.

For details about this step, see [Create a traffic classification policy for FabricPool](#).

Steps

1. Review the policy.
2. If you want to create this traffic classification policy, select **Create and continue**.

As soon as FabricPool begins tiering data to StorageGRID, you can go to the Traffic Classification Policies page to view network traffic metrics for this policy. Later, you can also add rules to limit other workloads and ensure that the FabricPool workload has most of the bandwidth.

3. Otherwise, select **Skip this step**.

Step 9 of 9: Review summary

The summary provides details about the items you configured, including the name of the load balancer, tenant, and bucket, the traffic classification policy, and the active ILM policy,

Steps

1. Review the summary.

2. Select **Finish**.

Next steps

After completing the FabricPool wizard, perform these additional steps.

Steps

1. Go to [Configure ONTAP System Manager](#) to enter the saved values and to complete the ONTAP side of the connection. You must add StorageGRID as a cloud tier, attach the cloud tier to a local tier to create a FabricPool, and set volume tiering policies.
2. Go to [Configure the DNS server](#) and ensure that the DNS includes a record to associate the StorageGRID server name (fully qualified domain name) to each StorageGRID IP address you will use.
3. Go to [Other best practices for StorageGRID and FabricPool](#) to learn the best practices for StorageGRID audit logs and other global configuration options.

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.