



Back up cloud-native Oracle databases

BlueXP backup and recovery

NetApp
June 11, 2024

Table of Contents

- Back up cloud-native Oracle databases 1
 - Quick start 1
 - Configure FSx for ONTAP 2
 - Configure Cloud Volumes ONTAP 3
 - Configure Azure NetApp Files 3
 - Install SnapCenter Plug-in for Oracle and add database hosts 4
- Back up cloud-native Oracle databases 10

Back up cloud-native Oracle databases

Quick start

Get started quickly by following these steps.

1

Verify support for your configuration

- Operating System:
 - RHEL 7.5 or later and 8.x
 - OL 7.5 or later and 8.x
 - SLES 15 SP4
- NetApp Cloud Storage:
 - Amazon FSx for NetApp ONTAP
 - Cloud Volumes ONTAP
 - Azure NetApp Files
- Storage layouts:
 - NFS v3 and v4.1 (including dNFS)
 - iSCSI with ASM (ASMFD, ASMLib and ASMUdev)



Azure NetApp Files does not support SAN environment.

- Database layouts: Oracle Standard and Oracle Enterprise Standalone (legacy and multitenant CDB and PDB)
- Database versions: 19c and 21c

2

Sign up to BlueXP

BlueXP is accessible from a web-based console. When you get started with BlueXP, your first step is to sign up using your existing NetApp Support Site credentials or by creating a NetApp cloud login. For information, refer to [Sign up to BlueXP](#).

3

Log into BlueXP

After you sign up to BlueXP, you can log in from the web-based console. For information, refer to [Log into BlueXP](#).

4

Manage your BlueXP account

You can administer your account by managing users, service accounts, workspaces, and Connectors. For information, refer to [Manage your BlueXP account](#).

Configure FSx for ONTAP

Using BlueXP you should create an FSx for ONTAP working environment to add and manage volumes and additional data services. You should also create a Connector in AWS that enables BlueXP to manage resources and processes within your public cloud environment.

Create FSx for ONTAP working environment

You should create the FSx for ONTAP working environments where your databases are hosted. For information, refer to [Get started with Amazon FSx for ONTAP](#) and [Create and manage an Amazon FSx for ONTAP working environment](#).

You can create the FSx for ONTAP working environment either using BlueXP or AWS. If you have created using AWS, then you should discover the FSx for ONTAP systems in BlueXP.

Create a Connector

An Account Admin needs to create a Connector in AWS that enables BlueXP to manage resources and processes within your public cloud environment.

For information, refer to [Creating a Connector in AWS from BlueXP](#).

- You should use the same connector to manage both FSx for ONTAP working environment and databases.
- If you have the FSx for ONTAP working environment and databases in the same virtual private cloud (VPC), you can deploy the connector in the same VPC.
- If you have the FSx for ONTAP working environment and databases in different VPCs:
 - If you have NAS (NFS) workloads configured on FSx for ONTAP, then you can create the connector on either of the VPCs.
 - If you have only SAN workloads configured and not planning to use any NAS (NFS) workloads, then you should create the connector in the VPC where the FSx for ONTAP system is created.



For using NAS (NFS) workloads, you should have transit gateway between the database VPC and Amazon VPC. The NFS IP address which is a floating IP address can be accessed from another VPC only through transit gateway. We cannot access the floating IP addresses by peering the VPCs.

After creating the Connector, click **Storage > Canvas > My Working Environments > Add Working Environment** and follow the prompts to add the working environment.

Ensure that there is connectivity from the Connector to the Oracle database hosts and FSx working environment. The Connector should be able to connect to the cluster management IP address of the FSx working environment.

- Add the working environment by clicking **Storage > Canvas > My Working Environments > Add Working Environment**.

Ensure that there is connectivity from the connector to the database hosts and FSx for ONTAP working environment. The connector should connect to the cluster management IP address of the FSx for ONTAP working environment.

- Copy the Connector ID by clicking **Connector > Manage Connectors** and selecting the Connector name.

Configure Cloud Volumes ONTAP

Using BlueXP you should create a Cloud Volumes ONTAP working environment to add and manage volumes and additional data services. You should also create a Connector for your cloud environment that enables BlueXP to manage resources and processes within your public cloud environment.

Create Cloud Volumes ONTAP working environment

You can discover and add existing Cloud Volumes ONTAP systems to BlueXP. For information, refer to [Adding existing Cloud Volumes ONTAP systems to BlueXP](#).

Create a Connector

You can get started with Cloud Volumes ONTAP for your cloud environment in a few steps. For more information, refer one of the following:

- [Quick start for Cloud Volumes ONTAP in AWS](#)
- [Quick start for Cloud Volumes ONTAP in Azure](#)
- [Quick start for Cloud Volumes ONTAP in Google Cloud](#)

You should use the same connector to manage both Cloud Volumes ONTAP working environment and databases.

- If you have the Cloud Volumes ONTAP working environment and databases in the same virtual private cloud (VPC) or VNet, you can deploy the connector in the same VPC or VNet.
- If you have the Cloud Volumes ONTAP working environment and databases in different VPCs or VNets, ensure that the VPCs or VNets are peered.

Configure Azure NetApp Files

Using BlueXP you should create a Azure NetApp Files working environment to add and manage volumes and additional data services. You should also create a Connector in Azure that enables BlueXP to manage resources and processes within your public cloud environment.

Create Azure NetApp Files working environment

You should create Azure NetApp Files working environments where your databases are hosted. For more information, refer to [Learn about Azure NetApp Files](#) and [Create an Azure NetApp Files working environment](#).

Create a connector

A BlueXP account admin should deploy a Connector in Azure that enables BlueXP to manage resources and processes within your public cloud environment.

For information, refer to [Create a Connector in Azure from BlueXP](#).

- Ensure that there is connectivity from the connector to the database hosts.
- If you have the Azure NetApp Files working environment and databases in the same Virtual Network (VNet), you can deploy the connector in the same VNet.
- If you have the Azure NetApp Files working environment and databases in different VNets and have NAS (NFS) workloads configured on Azure NetApp Files, then you can create the connector on either of the VNets.

After creating the connector, add the working environment by clicking **Storage > Canvas > My Working Environments > Add Working Environment**.

Install SnapCenter Plug-in for Oracle and add database hosts

You should install the SnapCenter Plug-in for Oracle on each of the Oracle database hosts, add the database hosts, and discover the databases on the host to assign policies and create backups.

- If SSH is enabled for the database host, you can install the plug-in using one of the methods:
 - Install the plug-in and add host from the UI using SSH option. [Learn more](#).
 - Install the plug-in using script and add host from the UI using manual option. [Learn more](#).
- If SSH is disabled, install the plug-in manually and add host from the UI using manual option. [Learn more](#).

Prerequisites

Before adding the host, you should ensure that the prerequisites are met.

- You should have created the working environment and the Connector.
- Ensure that the Connector has connectivity to the Oracle database hosts.

For information on how to resolve the connectivity issue, refer to [Failed to validate connectivity from BlueXP connector host to application database host](#).

When the connector is lost or if you have created a new connector, you should associate the connector with the existing application resources. For instructions to update the Connector, see [Update the Connector Details](#).

- Ensure that the BlueXP user has the “Account Admin” role.
- Ensure that non root (sudo) account is present on the application host for data protection operations.
- Ensure that either Java 11 (64-bit) Oracle Java or OpenJDK is installed on each of the Oracle database hosts and the JAVA_HOME variable is set appropriately.
- Ensure that the Connector has the communication enabled to the SSH port (default: 22) if SSH based installation is performed.
- Ensure that the Connector has the communication enabled to plug-in port (default: 8145) for the data protection operations to work.
- Ensure that you have the latest version of plug-in is installed. To upgrade the plug-in, refer to [Upgrade SnapCenter Plug-in for Oracle Database](#).

Add host from UI using SSH option

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**.

If you have already added a host and want to add another host, click **Applications > Manage Databases > Add** and then proceed with step 5.

2. Click **Discover Applications**.
3. Select **Cloud Native** and click **Next**.

A service account (*SnapCenter-account-`<accountid>`*) with *SnapCenter System* role is created to perform scheduled data protection operations for all the users in this account.

The service account (*SnapCenter-account-`<accountid>`*) is used to run the scheduled backup operations. You should never delete the service account.

You can view the service account by clicking **Account > Manage Account > Members**.

4. Select Oracle as the application type.
5. In the Host details page, perform the following:

- a. Select **Using SSH**.
- b. Specify the FQDN or IP address of the host where you want to install the plug-in.

Ensure that the Connector can communicate with the database host using the FQDN or IP address.

- c. Specify the non-root(sudo) user using which the plug-in package will be copied to the host.

Root user is not supported.

- d. Specify the SSH and plug-in port.

Default SSH port is 22 and the plug-in port is 8145.

You can close the SSH port on the application host after installing the plug-in. The SSH port is not required for any data protection operations.

- e. Select the Connector.
- f. (Optional) If key less authentication is not enabled between the Connector and the host, you should specify the SSH private key that will be used to communicate with the host.



The SSH private key is not stored anywhere in the application and is not used for any other operations.

- g. Click **Next**.
6. In the Configuration page, perform the following:
 - a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database.
 - b. Copy the text displayed in BlueXP UI.
 - c. Create the `/etc/sudoers.d/snapcenter` file on the Linux machine and paste the copied text.
 - d. In the BlueXP UI, select the checkbox and click **Next**.

7. Review the details and click **Discover Applications**.

- After the plug-in is installed, the discovery operation starts.
- After completing the discovery operation, all the databases on the host are displayed. If OS authentication is disabled for the database, click **Configure** to enable database authentication. For more information, refer to [Configure Oracle database credentials](#).
- Click **Settings** and select **Hosts** to view all the hosts.
- Click **Settings** and select **Policies** to view the pre-canned policies. Review the pre-canned policies and you can either edit them to meet your requirement or create a new policy.

Add host from UI using manual option and install the plug-in using script

Configure SSH key based authentication for the Oracle host non-root user account and perform the following steps to install the plug-in.

Before you begin

Ensure that the SSH connection to the Connector is enabled.

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. Click **Discover Applications**.
3. Select **Cloud Native** and click **Next**.

A service account (*SnapCenter-account-`<accountid>`*) with *SnapCenter System* role is created to perform scheduled data protection operations for all the users in this account.

The service account (*SnapCenter-account-`<accountid>`*) is used to run the scheduled backup operations. You should never delete the service account.

You can view the service account by clicking **Account > Manage Account > Members**.

4. Select Oracle as the application type.
5. In the Host details page, perform the following:
 - a. Select **Manual**.
 - b. Specify the FQDN or IP address of the host where the plug-in is installed.

Ensure that the Connector can communicate with the database host using the FQDN or IP address.

- c. Specify the plug-in port.

Default port is 8145.

- d. Specify the non-root (sudo) user using which the plug-in package will be copied to the host.
- e. Select the Connector.
- f. Select the check box to confirm that the plug-in is installed on the host.
- g. Click **Next**.

6. In the Configuration page, perform the following:
 - a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database.
 - b. Copy the text displayed in BlueXP UI.

c. Create the `/etc/sudoers.d/snapcenter` file on the Linux machine and paste the copied text.

d. In the BlueXP UI, select the checkbox and click **Next**.

7. Log into the Connector VM.

8. Install the plug-in using the script provided in the Connector.

```
sudo /var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh --host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port>
```

If you are using an older Connector, run the following command to install the plug-in.

```
sudo /var/lib/docker/volumes/cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh --host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port>
```

Name	Description	Mandatory	Default
plugin_host	Specifies the Oracle host	Yes	-
host_user_name	Specifies the SnapCenter user with SSH privileges on the Oracle host	Yes	-
host_ssh_key	Specifies the SSH key of the SnapCenter user and is used to connect to the Oracle host	Yes	-
plugin_port	Specifies the port used by the plug-in	No	8145
host_ssh_port	Specifies the SSH port on the Oracle host	No	22

For example:

- `sudo /var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh --host 10.0.1.1 --username snapcenter --sshkey /keys/netapp-ssh.ppk`

- `sudo /var/lib/docker/volumes/cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh --host 10.0.1.1 --username snapcenter --sshkey /keys/netapp-ssh.ppk`

9. In the BlueXP UI, review the details and click **Discover Applications**.

- After completing the discovery operation, all the databases on the host are displayed. If OS authentication is disabled for the database, click **Configure** to enable database authentication. For more information, refer to [Configure Oracle database credentials](#).

- Click **Settings** and select **Hosts** to view all the hosts.
- Click **Settings** and select **Policies** to view the pre-canned policies. Review the pre-canned policies and you can either edit them to meet your requirement or create a new policy.

Add host from UI using manual option and install the plug-in manually

If SSH key based authentication is not enabled on the Oracle database host, you should perform the following manual steps to install the plug-in and then add the host from UI using manual option.

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. Click **Discover Applications**.
3. Select **Cloud Native** and click **Next**.

A service account (*SnapCenter-account- \langle accountid \rangle*) with *SnapCenter System* role is created to perform scheduled data protection operations for all the users in this account.

The service account (*SnapCenter-account- \langle accountid \rangle*) is used to run the scheduled backup operations. You should never delete the service account.

You can view the service account by clicking **Account > Manage Account > Members**.

4. Select Oracle as the application type.
5. In the **Host details** page, perform the following:
 - a. Select **Manual**.
 - b. Specify the FQDN or IP address of the host where the plug-in is installed.

Ensure that using the FQDN or IP address, the Connector can communicate with the database host.
 - c. Specify the plug-in port.

Default port is 8145.
 - d. Specify the sudo non-root (sudo) user using which the plug-in package will be copied to the host.
 - e. Select the Connector.
 - f. Select the check box to confirm that the plug-in is installed on the host.
 - g. Click **Next**.
6. In the Configuration page, perform the following:
 - a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database.
 - b. Copy the text displayed in BlueXP UI.
 - c. Create the */etc/sudoers.d/snapcenter* file on the Linux machine and paste the copied text.
 - d. In the BlueXP UI, select the checkbox and click **Next**.

7. Log into the Connector VM.

8. Download the SnapCenter Linux host plug-in binary.

```
sudo docker exec -it cloudmanager_scs_cloud curl -X GET
'http://127.0.0.1/deploy/downloadLinuxPlugin'
```

The plug-in binary is available at: `cd /var/lib/docker/volumes/service-manager[1]-`

```
2_cloudmanager_scs_cloud_volume/_data/$(sudo docker ps|grep -Po "cloudmanager_scs_cloud:.*?"|sed -e 's/ *$//'|cut -f2 -d":")/sc-linux-host-plugin
```

9. Copy `snapcenter_linux_host_plugin_scs.bin` from the above path to `/home/<non root user (sudo)>/sc_netapp` path for each of the Oracle database hosts either using `scp` or other alternate methods.
10. Log into the Oracle database host using the non-root (`sudo`) account.
11. Change directory to `/home/<non root user>/sc_netapp/` and run the following command to enable execute permissions for the binary.

```
chmod +x snapcenter_linux_host_plugin_scs.bin
```
12. Install the Oracle plug-in as a sudo SnapCenter user.

```
./snapcenter_linux_host_plugin_scs.bin -i silent -DSPL_USER=<non-root>
```
13. Copy `certificate.pem` from `<base_mount_path>/client/certificate/` path of the Connector VM to `/var/opt/snapcenter/spl/etc/` on the plug-in host.
14. Navigate to `/var/opt/snapcenter/spl/etc` and execute the `keytool` command to import the `certificate.pem`.

```
keytool -import -alias agentcert -file certificate.pem -keystore keystore.jks -deststorepass snapcenter -noprompt
```
15. Restart SPL:

```
systemctl restart spl
```
16. Validate that the plug-in is reachable from the Connector by running the below command from the Connector.

```
docker exec -it cloudmanager_scs_cloud curl -ik https://<FQDN or IP of the plug-in host>:<plug-in port>/PluginService/Version --cert /config/client/certificate/certificate.pem --key /config/client/certificate/key.pem
```
17. In the BlueXP UI, review the details and click **Discover Applications**.
 - After completing the discovery operation, all the databases on the host are displayed. If OS authentication is disabled for the database, click **Configure** to enable database authentication. For more information, refer to [Configure Oracle database credentials](#).
 - Click **Settings** and select **Hosts** to view all the hosts.
 - Click **Settings** and select **Policies** to view the pre-canned policies. Review the pre-canned policies and you can either edit them to meet your requirement or create a new policy.

Configure Oracle database credentials

You should configure the database credentials that are used to perform data protection operations on Oracle databases.

Steps

1. If OS authentication is disabled for the database, click **Configure** to modify database authentication.
2. Specify the username, password, and the port details.

If the database is residing on ASM, you should also configure the ASM settings.

The Oracle user should have `sysdba` privileges and ASM user should have `sysasm` privileges.

3. Click **Configure**.

Upgrade SnapCenter Plug-in for Oracle Database

You should upgrade the SnapCenter Plug-in for Oracle to gain access to the latest new features and enhancements. You can upgrade from the BlueXP UI or using the command line.

Before you begin

- Ensure that there are no operations running on the host.

Steps

1. Click **Backup and recovery > Applications > Hosts**.
2. Verify if plug-in upgrade is available for any of the hosts by checking the Overall Status column.
3. Upgrade the plug-in from UI or using the command line.

Upgrade using UI	Upgrade using command line
<ol style="list-style-type: none">1. Click ... corresponding to the host and click Upgrade Plug-in.2. In the Configuration page, perform the following:<ol style="list-style-type: none">a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database.b. Copy the text displayed in BlueXP UI.c. Edit the <code>/etc/sudoers.d/snapcenter</code> file on the Linux machine and paste the copied text.d. In the BlueXP UI, select the checkbox and click Upgrade.	<ol style="list-style-type: none">1. Log in to Connector VM.2. Run the following script.<pre>sudo /var/lib/docker/volumes/service- manager- 2_cloudmanager_scs_cloud_volume/_da ta/scripts/linux_plugin_copy_and_in stall.sh --host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port> --upgrade</pre><p>If you are using an older Connector, run the following command to upgrade the plug-in.</p><pre>sudo /var/lib/docker/volumes/cloudmanage r_scs_cloud_volume/_data/scripts/li nux_plugin_copy_and_install.sh --host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port> --upgrade</pre>

Back up cloud-native Oracle databases

You can create scheduled or on-demand backups by assigning a pre-canned policy or the policy that you created.

You can also catalog the Oracle database backups using Oracle Recovery Manager (RMAN) if you have enabled cataloging while creating a policy. The (RMAN) cataloging is supported only for the databases that are on Azure NetApp Files. The cataloged backups can be used later for block-level restore or tablespace point-in-time recovery operations. The database must be in mounted or higher state for cataloging.

Create policy to protect Oracle database

You can create policies if you do not want to edit the pre-canned policies.

Steps

1. In the Applications page, from the Settings drop-down list, select **Policies**.
2. Click **Create policy**.
3. Specify a policy name.
4. (Optional) Edit the format of the backup name.
5. Specify the schedule and retention details.
6. If you have selected *daily* and *weekly* as the schedule and you want to enable RMAN cataloging, select **Catalog backup with Oracle Recovery Manager (RMAN)**.
7. (Optional) Enter the post-script path and timeout value for post-script that will be executed after the successful backup such as copying the snapshot to secondary storage.

Optionally, you can also specify the arguments.

You should keep the post-scripts in the path `/var/opt/snapcenter/spl/scripts`.

The post script supports a set of environment variables.

Environmental Variable	Description
SC_ORACLE_SID	Specifies the SID of the Oracle database.
SC_HOST	Specifies the hostname of the database
SC_BACKUP_NAME	Specifies the name of the backup. The data backup name and the log backup name are concatenated using delimiters.
SC_BACKUP_POLICY_NAME	Specifies the name of the policy used to create the backup.
SC_PRIMARY_DATA_VOLUME_FULL_PATH	Specifies the data volume paths concatenated using "," as delimiter. For Azure NetApp Files volumes, the information is concatenated using "/" _/ _subscriptions/{subscription_id}/resourceGroups/{resource_group}/providers/{provider}/netAppAccounts/{anfaccount}/capacityPools/{capacity_pool}/volumes/{volumename}_

Environmental Variable	Description
SC_PRIMARY_ARCHIVELOGS_VOLUME_FULL_PATH	Specifies the archive log volume paths concatenated using "," as delimiter. For Azure NetApp Files volumes, the information in concatenated using "/" _ /subscriptions/{subscription_id}/resourceGroups/{resource_group}/providers/{provider}/netAppAccounts/{anfaccount}/capacityPools/{capacity_pool}/volumes/{volumename}_

8. Click **Create**.


Configure RMAN catalog repository

You can configure the recovery catalog database as the RMAN catalog repository. If you do not configure the repository, by default, the Control file of the target database becomes the RMAN catalog repository.

Before you begin

You should manually register the target database with RMAN catalog database.

Steps

1. In the Applications page, click **...** > **View Details**.
2. In the Database details section, click  to configure the RMAN catalog repository.
3. Specify the credentials to catalog backups with RMAN and the Transparent Network Substrate (TNS) name of catalog recovery database.
4. Click **Configure**.

Create a backup of the Oracle Database

You can assign a pre-canned policy or create a policy and then assign it to the database. Once the policy is assigned, the backups are created as per the schedule defined in the policy.



When creating ASM diskgroups on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP, ensure that there are no common volumes across diskgroups. Each diskgroup should have dedicated volumes.

Steps

1. In the Applications page, if the database is not protected using any policy, click **Assign Policy**.

If the database is protected using one or more policies, you can assign more policies by clicking **...** > **Assign Policy**.

2. Select the policy and click **Assign**.

The backups will be created as per the schedule defined in the policy. If you have enabled RMAN catalog in the policy, the backup at the end of the workflow will launch the cataloging operation as a separate job. The progress of cataloging can be seen from the Job Monitor. Upon successful cataloging, **Backup Details**

will show the status of the catalog for each backup.



The service account (*SnapCenter-account- \langle account_id \rangle*) is used to run the scheduled backup operations.

Create on-demand backup of the Oracle database

After assigning the policy, you can create an on-demand backup of the application.

Steps

1. In the Applications page, click **...** corresponding to the application and click **On-Demand Backup**.
2. If multiple policies are assigned to the application, select the policy, retention tier, and then click **Create Backup**.

If you have enabled RMAN catalog in the policy, the backup at the end of the workflow will launch the cataloging operation as a separate job. The progress of cataloging can be seen from the Job Monitor. Upon successful cataloging, **Backup Details** will show the status of the catalog for each backup.

Limitations

- Does not support consistency group Snapshots for Oracle databases residing on Multiple ASM disk groups with overlap of FSx volumes
- If your Oracle databases are on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP and are configured on ASM, ensure your SVM names are unique across the FSx systems. If you have same SVM name across FSx systems, back up of Oracle databases residing on those SVMs are not supported.
- After restoring a large database (250 GB or more), if you perform a full online backup on the same database the operation might fail with the following error:
failed with status code 500, error

```
{\"error\":{\"code\":\"app_internal_error\",\"message\":\"Failed to create snapshot. Reason: Snapshot operation not allowed due to clones backed by snapshots. Try again after sometime.
```

For information on how to fix this issue, refer to: [Snapshot operation not allowed due to clones backed by snapshots](#).

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.