



# Back up and restore cloud native applications data

## Cloud Backup

NetApp  
December 07, 2022

This PDF was generated from <https://docs.netapp.com/us-en/cloud-manager-backup-restore/concept-protect-cloud-app-data-to-cloud.html> on December 07, 2022. Always check docs.netapp.com for the latest.

# Table of Contents

- Back up and restore cloud native applications data ..... 1
  - Protect your cloud native applications data ..... 1
  - Prerequisites ..... 4
  - Back up cloud native applications data ..... 6
  - Manage protection of cloud native application data ..... 8
  - Restore cloud native application data ..... 12

# Back up and restore cloud native applications data

## Protect your cloud native applications data

Cloud Backup for Applications is a SaaS based service that provides data protection capabilities for applications running on NetApp Cloud Storage. Cloud Backup for Applications enabled within NetApp BlueXP (formerly Cloud Manager) offers efficient, application consistent, policy-based backup, and restore of Oracle databases residing on Amazon FSx for NetApp ONTAP.

### Architecture

The Cloud Backup for Applications architecture includes the following components.

- Cloud Backup for Applications is a set of data protection services hosted as a SaaS service by NetApp and is based on the BlueXP SaaS platform.

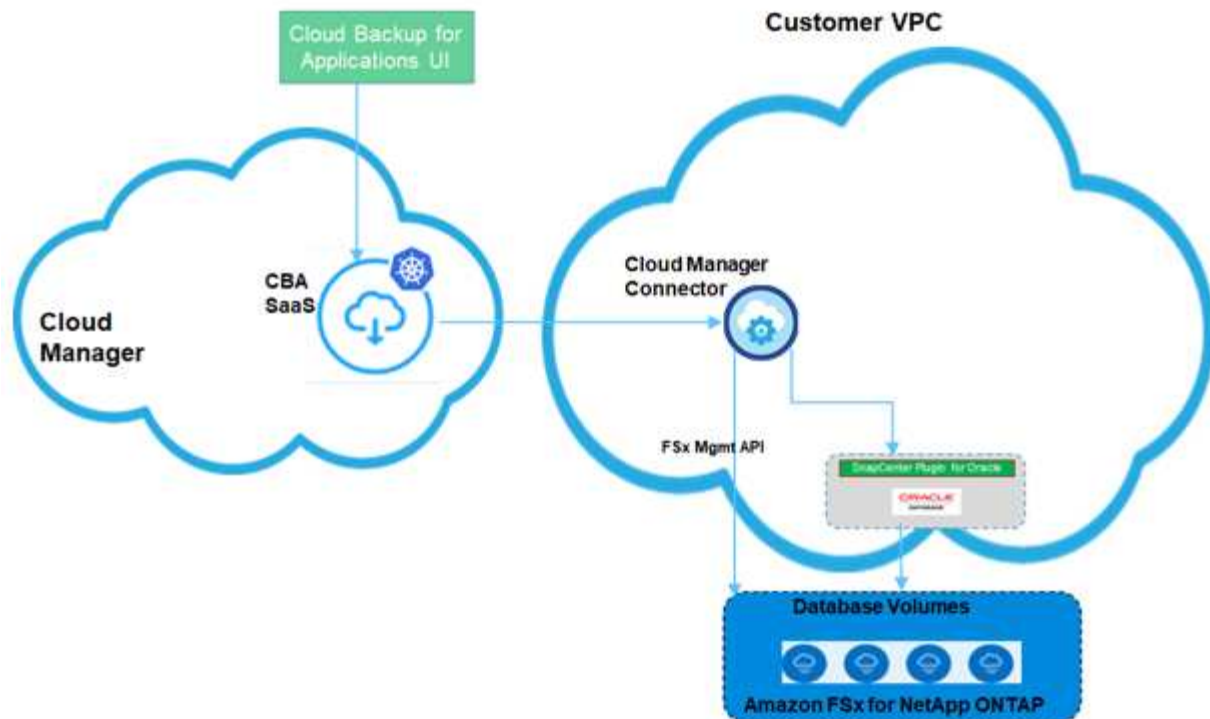
It orchestrates the data protection workflows for applications residing on NetApp Cloud Storage.

- Cloud Backup for Applications UI is integrated with BlueXP UI.

The Cloud Backup for Applications UI offers multiple storage and data management capabilities.

- BlueXP Connector is a component from BlueXP that runs in your cloud network and interacts with Amazon FSx storage file systems and SnapCenter Plug-in for Oracle running on Oracle database hosts.
- SnapCenter Plug-in for Oracle is a component that runs on each Oracle database hosts and interacts with the Oracle databases running on the host while performing data protection operations.

The following image shows each component and the connections that you need to prepare between them:



For any user-initiated request, the Cloud Backup for Applications UI communicates with the BlueXP SaaS which upon validating the request processes the same. If the request is to run a workflow such as a backup or restore, the SaaS service initiates the workflow and where required, forwards the call to the BlueXP Connector. The Connector then communicates with Amazon FSx for NetApp ONTAP and SnapCenter Plug-in for Oracle as part of running the workflow tasks.

The Connector can be deployed in the same VPC as that of the Oracle databases, or in a different one. If the Connector and Oracle databases are on different network, you should establish a network connectivity between them.



Cloud Backup for Applications infrastructure is resilient to availability zone failures within a region. It now supports regional failures by failing over to a new region and this failover involves a downtime of around 2 hours.

## Supported configurations

- Operating System:
  - RHEL 7.5 or later and 8.x
  - OL 7.5 or later and 8.x
- Storage System: Amazon FSx for ONTAP
- Storage layouts: NFS v3 and v4.1 (dNFS is supported) and iSCSI with ASM (ASMFD, ASMLib and ASMUdev)
- Applications: Oracle Standard and Oracle Enterprise – Standalone (legacy and multitenant – CDB and PDB)
- Oracle versions: 12cR2, 18c, and 19c

## Features

- Auto-discovery of Oracle databases
- Backing up Oracle databases residing on Amazon FSx for NetApp ONTAP
  - Full (data + control + archive log files) backup
  - On-demand backup
  - Scheduled backup based on the system-defined or custom policies

You can specify different scheduling frequencies such as hourly, daily, weekly, and monthly in the policy.

- Retaining backups based on the policy
- Restoring complete Oracle database (data files + control file) from the specified backup
- Restoring data files only and control files only from the specified backup
- Recovering Oracle database with until SCN, until time, all available logs, and no recovery options
- Monitoring backups and other jobs
- Displaying the protection summary on the dashboard
- Sending alerts through email

## Limitations

- Does not support Oracle versions 11g and 21c
- Does not support mount, clone, catalog, and verification operations on backups
- Does not support Oracle on RAC and Data Guard
- Backup limitations:
  - Does not support online data or log only backups
  - Does not support offline backups
  - Does not support backing up of Oracle database residing on recursive mount points
  - 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 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.
- Restore limitations:
  - Does not support granular restores, for example restoring of tablespaces and PDBs
  - Supports only in-place restore of Oracle databases on NAS and SAN Layouts
  - Does not support restore of control file only or data files + control file of Oracle databases on SAN layouts
  - In SAN layout, restore operation fails if SnapCenter Plug-in for Oracle finds any foreign files other than Oracle data files on the ASM diskgroup. The foreign files could be one or more of the following types:
    - Parameter
    - Password

- archive log
- online log
- ASM parameter file.

You should select Force in-place restore checkbox to override the foreign files of type parameter, password, and archive log.



If there are other types of foreign files, restore operation fails and the database cannot be recovered. If you have other type of foreign files, you should delete or move them to a different location before performing restore operation.

The failure message due to the presence of foreign files are not displayed in the job page in UI due to a known issue. Check the connector logs if there is a failure during SAN pre-restore stage to know the cause of the issue.

## Prerequisites

You should have access to BlueXP, created a BlueXP account, created the working environment and a Connector, and deployed the SnapCenter Plug-in for Oracle.

### Access BlueXP

You should [log into BlueXP](#), and then set up a [NetApp account](#).

### Create FSx for ONTAP working environment

You should create the Amazon 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 NetApp FSx 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 deploy 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 working environment and Oracle databases.
- If you have the FSx working environment and Oracle databases in the same VPC, you can deploy the connector in the same VPC.
- If you have the FSx working environment and Oracle databases in different VPCs:
  - If you have NAS (NFS) workloads configured on FSx, 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 system is created.



For using NAS (NFS) workloads, you should have transit gateway between the Oracle database VPC and FSx 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.



After creating the Connector, click **Connector > Manage Connectors**; select the Connector name and copy the Connector ID.

## Deploy SnapCenter Plug-in for Oracle

You should deploy the SnapCenter Plug-in for Oracle on each of the Oracle database hosts. Depending on whether the Oracle host has SSH key based authentication enabled, you can follow one of the methods to deploy the plug-in.



Ensure that JAVA 8 is installed on each of the Oracle database hosts and the JAVA\_HOME variable is set appropriately.

### Plug-in deployment using SSH key based authentication

If SSH key based authentication is enabled on the Oracle host, you can perform the following steps to deploy the plug-in. Before performing the steps, ensure that the SSH connection to the Connector is enabled.

1. Log into the Connector VM as non root user.

2. Obtain the base mount path.

```
sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs sudo
docker volume inspect | grep Mountpoint
```

3. Deploy the plug-in.

```
sudo <base_mount_path>/scripts/oracle_plugin_copy_and_install.sh --host
<host_name> --sshkey <ssh_key_file> --username <user_name> --port <ssh_port>
--pluginport <plugin_port> --installdir <install_dir>
```

- `host_name` is the name of the Oracle host and this is a mandatory parameter.
- `ssh_key_file` is SSH key used to connect to the Oracle host and this is a mandatory parameter.
- `user_name`: User with SSH privileges on the Oracle host and this is an optional parameter. Default value is `ec2-user`.
- `ssh_port`: SSH port on the Oracle host and this is an optional parameter. Default value is 22
- `plugin_port`: Port used by the plug-in and this is an optional parameter. Default value is 8145
- `install_dir`: Directory where the plug-in will be deployed and this is an optional parameter. Default value is `/opt`.

For example: `sudo /var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/scripts/oracle_plugin_copy_and_install.sh`

```
--host xxx.xx.x.x --sshkey /keys/netapp-ssh.ppk
```

## Manual deployment of the plug-in

If SSH key based authentication is not enabled on the Oracle host, you should perform the following manual steps to deploy the plug-in.

1. Log into the Connector VM.

2. 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'
```

3. Obtain the base mount path.

```
sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs sudo  
docker volume inspect | grep Mountpoint
```

4. Obtain the binary path of the downloaded plug-in.

```
sudo ls <base_mount_path> $(sudo docker ps|grep -Po  
"cloudmanager_scs_cloud:.*? "|sed -e 's/ *$//'|cut -f2 -d":")/sc-linux-host  
-plugin/snapcenter_linux_host_plugin_scs.bin
```

5. Copy *snapcenter\_linux\_host\_plugin\_scs.bin* to each of the Oracle database hosts either using scp or other alternate methods.

6. On the Oracle database host, run the following command to enable execute permissions for the binary.

```
chmod +x snapcenter_linux_host_plugin_scs.bin
```

7. Deploy the Oracle plug-in as a root user.

```
./snapcenter_linux_host_plugin_scs.bin -i silent
```

8. Copy *certificate.p12* from *<base\_mount\_path>/client/certificate/* path of the Connector VM to */var/opt/snapcenter/spl/etc/* on the plug-in host.

a. Navigate to */var/opt/snapcenter/spl/etc* and execute the keytool command to import the certificate.

```
keytool -v -importkeystore -srckeystore certificate.p12 -srcstoretype PKCS12  
-destkeystore keystore.jks -deststoretype JKS -srcstorepass snapcenter  
-deststorepass snapcenter -srccalias agentcert -destalias agentcert -noprompt
```

b. Restart SPL: `systemctl restart spl`

## Back up cloud native applications data

### Discover the applications

You should discover the databases on the host to assign policies and create backups.

#### What you will need

- You should have created the FSx for ONTAP working environment and the Connector.
- Ensure that the Connector has connectivity to the FSx for ONTAP working environment and Oracle database hosts.
- Ensure that the BlueXP user has the “Account Admin” role.
- You should have deployed the SnapCenter Plug-in for Oracle. [Learn more.](#)



## Steps

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

A service account with *SnapCenter System* role is created to perform scheduled data protection operations for all the users in this account.

- Click **Account > Manage Account > Members** to view the service account.



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

4. In the Specify Host Details page, enter the details of the Oracle database host, select the check box to confirm that the plug-in is installed on the host, and click **Discover**.
  - Displays all the databases on the host. If OS authentication is disabled for the database, you should configure database authentication by clicking **Configure**. For more information, refer to [Configure Oracle database credentials](#).
  - Click **Manage Application**, select **Add** to add a new host, **Refresh** to discover new databases, or **Remove** to remove a database host.
  - Click **Settings** and select **Policies** to view the pre-canned policies. Review the pre-canned policies and if you want you can either edit them to meet your requirement or create a new policy.

## Configure Oracle database credentials

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

### Steps

1. If OS authentication is disabled for the database, you should configure database authentication by clicking **Configure**.
2. Specify the username, password, and the port details either in the Database Settings or ASM Settings section.

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

3. Click **Configure**.

## Create policy

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. Click **Create**.

## Back up the cloud native application data

You can either 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. You can also create an on-demand backup.



When creating ASM diskgroups for Oracle, ensure that there are no common volumes across diskgroups. Each diskgroup needs to 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.



The service account (*SnapCenter-account-`<account_id>`*) is used for running the scheduled backup operations.

### Create on-demand backups

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 value, and then click **Create Backup**.

### Find more information

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](#).

## Manage protection of cloud native application data

## Monitor jobs

You can monitor the status of the jobs that have been initiated in your working environments. This allows you to see the jobs that have completed successfully, those currently in progress, and those that have failed so you can diagnose and fix any problems.

You can view a list of all the operations and their status. Each operation, or job, has a unique ID and a status. The status can be:

- Successful
- In Progress
- Queued
- Warning
- Failed

### Steps

1. Click **Backup and recovery**.
2. Click **Job Monitoring**

You can click the name of a job to view details corresponding to that operation. If you are looking for specific job, you can:

- use the time selector at the top of the page to view jobs for a certain time range
- enter a part of the job name in the Search field
- sort the results by using the filter in each column heading

## View backup details

You can view total number of backups created, policies used for creating backups, database version, and agent ID.

1. Click **Backup and recovery > Applications**.
2. Click **...** corresponding to the application and click **View Details**.



The agent ID is associated to the Connector. If a Connector that was used during registering the Oracle database host no longer exists, the subsequent backups of that application will fail because the agent ID of the new Connector is different. You should run the **connector-update** API to modify the agent ID.

## Update the Connector Details

If a Connector that was used during registering the Oracle database host no longer exists or is corrupted in AWS, you should deploy a new connector. After deploying the new connector, you should run the **connector-update** API to update the Connector details.

```

curl --location --request PATCH
'https://snapcenter.cloudmanager.cloud.netapp.com/api/oracle/databases/con
nector-update' \
--header 'x-account-id: <CM account-id>' \
--header 'x-agent-id: <connector Agent ID >' \
--header 'Authorization: Bearer token' \
--header 'Content-Type: application/json' \
--data-raw '{
"old_connector_id": "Old connector id that no longer exist",
"new_connector_id": "New connector Id"
}'

```

After updating the Connector details, you should connect to the Oracle database host and perform the following steps:

1. Obtain the plug-in information running on the Oracle database host.  
`rpm -qi netapp-snapcenter-plugin-oracle`
2. Uninstall the plug-in.  
`sudo /opt/NetApp/snapcenter/spl/installation/plugins/uninstall`
3. Verify that the plug-in is uninstalled successfully.  
`rpm -qi netapp-snapcenter-plugin-oracle`

After uninstalling the plug-in, you can deploy the plug-in. [Learn more](#).

## Configure CA signed certificate

You can configure CA signed certificate if you want to include additional security to your environment.

### Configure CA signed certificate for client certificate authentication

The connector uses a self-signed certificate to communicate with plug-in. The self-signed certificate is imported to the keystore by the installation script. You can perform the following steps to replace the self-signed certificate with CA signed certificate.

#### What you will need

You can run the following command to get the `<base_mount_path>`:

```

sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs sudo
docker volume inspect | grep Mountpoint

```

#### Steps

1. Login to Connector.
2. Delete all the existing files located at `<base_mount_path>/client/certificate` in the Connector virtual machine.
3. Copy the CA signed certificate and key file to the `<base_mount_path>/client/certificate` in the Connector virtual machine.

The file name should be `certificate.pem` and `key.pem`. The `certificate.pem` should have the entire chain of the certificates like intermediate CA and root CA.

4. Create the PKCS12 format of the certificate with the name `certificate.p12` and keep at `<base_mount_path>/client/certificate`.
5. Copy the `certificate.p12` and certificates for all the intermediate ca and root ca to the plug-in host at `/var/opt/snapcenter/spl/etc/`.
6. Log in to the plug-in host.
7. Navigate to `/var/opt/snapcenter/spl/etc` and run the `keytool` command to import the `certificate.p12` file.  

```
keytool -v -importkeystore -srckeystore certificate.p12 -srcstoretype PKCS12 -destkeystore keystore.jks -deststoretype JKS -srcstorepass snapcenter -deststorepass snapcenter -srcalias agentcert -destalias agentcert -noprompt
```
8. Import the root CA and intermediate certificates.  

```
keytool -import -trustcacerts -keystore keystore.jks -storepass snapcenter -alias trustedca -file <certificate.crt>
```



The `certfile.crt` refers to the certificates of root CA as well as intermediate CA.

9. Restart SPL: `systemctl restart spl`

## Configure CA signed certificate for server certificate of plug-in

The CA certificate should have the exact name of the Oracle plug-in host with which the Connector virtual machine communicates.

### What you will need

You can run the following command to get the `<base_mount_path>`:

```
sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs sudo docker volume inspect | grep Mountpoint
```

### Steps

1. Perform the following steps on the plug-in host:
  - a. Navigate to the folder containing the SPL's keystore `/var/opt/snapcenter/spl/etc`.
  - b. Create the PKCS12 format of the certificate having both certificate and key with alias `splkeystore`.
  - c. Add the CA certificate.  

```
keytool -importkeystore -srckeystore <CertificatePathToImport> -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS -srcalias splkeystore -destalias splkeystore -noprompt
```
  - d. Verify the certificates.  

```
keytool -list -v -keystore keystore.jks
```
  - e. Restart SPL: `systemctl restart spl`
2. Perform the following steps on the Connector:
  - a. Log in to the Connector as non-root user.
  - b. Copy the entire chain of CA certificates to the persistent volume located at `<base_mount_path>/server`.

Create the server folder if it does not exist.

- c. Connect to the `cloudmanager_scs_cloud` and modify the **enableCACert** in `config.yml` to **true**.

```
sudo docker exec -t cloudmanager_scs_cloud sed -i 's/enableCACert:
false/enableCACert: true/g' /opt/netapp/cloudmanager-scs-
cloud/config/config.yml
```

- d. Restart `cloudmanager_scs_cloud` container.

```
sudo docker restart cloudmanager_scs_cloud
```

## Access REST APIs

The REST APIs to protect the applications to cloud is available [here](#).

You should obtain the user token with federated authentication to access the REST APIs. For information to obtain the user token, refer to [Create a user token with federated authentication](#).

## Restore cloud native application data

In the event of data loss, you can restore the data files, control files, or both and then recover the database.

### Steps

1. Click **...** corresponding to the database that you want to restore and click **View Details**.
2. Click **...** corresponding to the data backup that you want to use for restoring and click **Restore**.
3. In the Restore Scope section, perform the following actions:

If you...	Do this...
Want to restore only the data files	Select <b>All Data Files</b> .
Want to restore only the control files	Select <b>Control Files</b>
Want to restore both data files and control files	Select <b>All Data Files</b> and <b>Control Files</b> .



Restore of datafiles with control files or only control files are not supported for iSCSI on ASM layout.

You can also select **Force in-place restore** checkbox.

The **Force in-place restore** option overrides the spfile, password file, and archive log files from the data files diskgroup. You should use the latest backup when **Force in-place restore** option is selected.

4. In the Recovery Scope section, perform the following actions:

If you...	Do this...
Want to recover to the last transaction	Select <b>All Logs</b> .

If you...	Do this...
Want to recover to a specific System Change Number (SCN)	Select <b>Until System Change Number</b> and specify the SCN.
Want to recover to a specific date and time	Select <b>Date and Time</b> .
Do not want to recover	Select <b>No recovery</b> .

For the selected recovery scope, in the **Archive Log Files Locations** field you can optionally specify the location that contains the archive logs required for recovery.

Select the check box if you want to open the database in READ-WRITE mode after recovery.

5. Click **Next** and review the details.
6. Click **Restore**.

## Copyright information

Copyright © 2022 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.