



Protect Unix file systems

SnapCenter software

NetApp

January 09, 2026

This PDF was generated from https://docs.netapp.com/us-en/snapcenter/protect-scu/concept_overview_snapcenter_plug_in_for_UNIX_file_systems.html on January 09, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Protect Unix file systems	1
What you can do with the SnapCenter Plug-in for Unix file systems	1
Supported configurations	1
Limitations	2
Features	2
Install SnapCenter Plug-in for Unix file systems	2
Prerequisites for adding hosts and installing Plug-ins Package for Linux	2
Add hosts and install Plug-ins Package for Linux using GUI	3
Configure the SnapCenter Plug-in Loader service	6
Configure CA certificate with SnapCenter Plug-in Loader (SPL) service on Linux host	9
Enable CA Certificates for plug-ins	12
Install SnapCenter Plug-in for VMware vSphere	12
Deploy CA certificate	12
Configure the CRL file	13
Prepare for protecting Unix file systems	13
Back up Unix file systems	13
Discover the UNIX file systems available for backup	13
Create backup policies for Unix file systems	14
Create resource groups and attach policies for Unix file systems	16
Create resource groups and enable secondary protection for Unix file systems on ASA r2 systems	18
Back up Unix file systems	21
Back up Unix file systems resource groups	22
Monitor Unix file systems backup	22
View protected Unix file systems in the Topology page	24
Restore and recover Unix file systems	26
Restore Unix file systems	26
Monitor Unix file systems restore operations	27
Clone Unix file systems	28
Clone Unix file system backup	28
Split a clone	29
Monitor Unix file systems clone operations	31

Protect Unix file systems

What you can do with the SnapCenter Plug-in for Unix file systems

When the Plug-in for Unix file systems is installed in your environment, you can use SnapCenter to back up, restore, and clone Unix file systems. You can also perform tasks supporting those operations.

- Discover resources
- Back up Unix file systems
- Schedule backup operations
- Restore file system backups
- Clone file system backups
- Monitor backup, restore, and clone operations

Supported configurations

Item	Supported configuration
Environments	<ul style="list-style-type: none">• Physical server• Virtual server <p>vVol datastores on both NFS and SAN. vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.</p>
Operating systems	<ul style="list-style-type: none">• Red Hat Enterprise Linux• Oracle Linux• SUSE Linux Enterprise Server (SLES)
File systems	<ul style="list-style-type: none">• SAN:<ul style="list-style-type: none">◦ Both LVM and non LVM based file systems◦ LVM over VMDK ext3, ext4, and xfs• NFS: NFS v3, NFS v4.x
Protocols	<ul style="list-style-type: none">• FC• FCoE• iSCSI• NFS
Multipath	yes

Limitations

- Mix of RDMs and virtual disks in a volume group is not supported.
- File level restore is not supported.

However, you can manually perform file level restore by cloning the backup and then copying the files manually.

- Mix of file systems spread across VMDKs coming from both NFS and VMFS datastore is not supported.
- NVMe is not supported.
- Provisioning is not supported.

Features

- Enables the Plug-in for Oracle Database to perform data protection operations on Oracle databases by handling the underlying host storage stack on Linux or AIX systems
- Supports Network File System (NFS) and storage area network (SAN) protocols on a storage system that is running ONTAP.
- For Linux systems, Oracle databases on VMDK and RDM LUNs is supported when you deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.
- Supports Mount Guard for AIX on SAN filesystems and LVM layout.
- Supports Enhanced Journaled File System (JFS2) with inline logging on SAN filesystems and LVM layout for AIX systems only.

SAN native devices, filesystems, and LVM layouts built on SAN devices are supported.

- Automates application-aware backup, restore, and clone operations for UNIX file systems in your SnapCenter environment

Install SnapCenter Plug-in for Unix file systems

Prerequisites for adding hosts and installing Plug-ins Package for Linux

Before you add a host and install the plug-ins package for Linux, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You can either use the password-based authentication for the root or non-root user or SSH key based authentication.

SnapCenter Plug-in for Unix File Systems can be installed by a non-root user. However, you should configure the sudo privileges for the non-root user to install and start the plug-in process. After installing the plug-in, the processes will be running as an effective non-root user.

- Create credentials with authentication mode as Linux for the install user.
- You must have installed Java 11 on your Linux host.



Ensure that you have installed only the certified edition of JAVA 11 on the Linux host.

For information to download JAVA, see: [Java Downloads for All Operating Systems](#)

- You should have **bash** as the default shell for plug-in installation.

Linux Host requirements

You should ensure that the host meets the requirements before installing the SnapCenter Plug-ins Package for Linux.

Item	Requirements
Operating systems	<ul style="list-style-type: none">• Red Hat Enterprise Linux• Oracle Linux• SUSE Linux Enterprise Server (SLES)
Minimum RAM for the SnapCenter plug-in on host	2 GB
Minimum install and log space for the SnapCenter plug-in on host	2 GB  You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 Oracle Java and OpenJDK  Ensure that you have installed only the certified edition of JAVA 11 on the Linux host. If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path.

For the latest information about supported versions, see the [NetApp Interoperability Matrix Tool](#).

Add hosts and install Plug-ins Package for Linux using GUI

You can use the Add Host page to add hosts, and then install the SnapCenter Plug-ins Package for Linux. The plug-ins are automatically installed on the remote hosts.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:

For this field...	Do this...
Host Type	Select Linux as the host type.
Host name	<p>Enter the fully qualified domain name (FQDN) or the IP address of the host.</p> <p>SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN.</p> <p>If you are adding a host by using SnapCenter and the host is part of a subdomain, you must provide the FQDN.</p>
Credentials	<p>Either select the credential name that you created or create new credentials.</p> <p>The credential must have administrative rights on the remote host. For details, see the information about creating credentials.</p> <p>You can view details about the credentials by positioning the cursor over the credential name that you specified.</p> <p> The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.</p>

5. In the Select Plug-ins to Install section, select **Unix File Systems**.
6. (Optional) Click **More Options**.

For this field...	Do this...
Port	<p>Either retain the default port number or specify the port number.</p> <p>The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port.</p> <p> If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.</p>
Installation Path	<p>The default path is <code>/opt/NetApp/snapcenter</code>.</p> <p>You can optionally customize the path. If you use the custom path, ensure that the default content of the sudoers is updated with the custom path.</p>
Skip optional preinstall checks	<p>Select this check box if you have already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.</p>

7. Click **Submit**.

If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in.



The precheck script does not validate the plug-in port firewall status if it is specified in the firewall reject rules.

Appropriate error or warning messages are displayed if the minimum requirements are not met. If the error is related to disk space or RAM, you can update the web.config file located at `C:\Program Files\NetApp\SnapCenter WebApp` to modify the default values. If the error is related to other parameters, you should fix the issue.



In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. Verify the fingerprint, and then click **Confirm and Submit**.



SnapCenter does not support ECDSA algorithm.



Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.

The installation-specific log files are located at `/custom_location/snapcenter/logs`.

Result

All the file systems mounted on the host are automatically discovered and displayed under the Resources Page. If nothing is displayed, click **Refresh Resources**.

Monitor installation status

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
 - a. Click **Filter**.
 - b. Optional: Specify the start and end date.
 - c. From the Type drop-down menu, select **Plug-in installation**.
 - d. From the Status drop-down menu, select the installation status.
 - e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure the SnapCenter Plug-in Loader service

The SnapCenter Plug-in Loader service loads the plug-in package for Linux to interact with the SnapCenter Server. The SnapCenter Plug-in Loader service is installed when you install the SnapCenter Plug-ins Package for Linux.

About this task

After installing the SnapCenter Plug-ins Package for Linux, the SnapCenter Plug-in Loader service starts automatically. If the SnapCenter Plug-in Loader service fails to start automatically, you should:

- Ensure that the directory where the plug-in is operating is not deleted
- Increase the memory space allotted to the Java Virtual Machine

The `spl.properties` file, which is located at `/custom_location/NetApp/snapcenter/spl/etc/`, contains the following parameters. Default values are assigned to these parameters.

Parameter name	Description
LOG_LEVEL	<p>Displays the log levels that are supported. The possible values are TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.</p>
SPL_PROTOCOL	<p>Displays the protocol that is supported by SnapCenter Plug-in Loader. Only the HTTPS protocol is supported. You can add the value if the default value is missing.</p>
SNAPCENTER_SERVER_PROTOCOL	<p>Displays the protocol that is supported by SnapCenter Server. Only the HTTPS protocol is supported. You can add the value if the default value is missing.</p>
SKIP_JAVAHOME_UPDATE	<p>By default, the SPL service detects the java path and update <code>JAVA_HOME</code> parameter. Therefore the default value is set to FALSE. You can set to TRUE if you want to disable the default behavior and manually fix the java path.</p>
SPL_KEYSTORE_PASS	<p>Displays the password of the keystore file. You can change this value only if you change the password or create a new keystore file.</p>
SPL_PORT	<p>Displays the port number on which the SnapCenter Plug-in Loader service is running. You can add the value if the default value is missing.</p> <p> You should not change the value after installing the plug-ins.</p>
SNAPCENTER_SERVER_HOST	Displays the IP address or host name of the SnapCenter Server.
SPL_KEYSTORE_PATH	Displays the absolute path of the keystore file.
SNAPCENTER_SERVER_PORT	Displays the port number on which the SnapCenter Server is running.

Parameter name	Description
LOGS_MAX_COUNT	<p>Displays the number of SnapCenter Plug-in Loader log files that are retained in the <code>/custom_location/snapcenter/spl/logs</code> folder.</p> <p>The default value is set to 5000. If the count is more than the specified value, then the last 5000 modified files are retained. The check for the number of files is done automatically every 24 hours from when SnapCenter Plug-in Loader service is started.</p> <p> If you manually delete the <code>spl.properties</code> file, then the number of files to be retained is set to 9999.</p>
JAVA_HOME	<p>Displays the absolute directory path of the <code>JAVA_HOME</code> which is used to start SPL service.</p> <p>This path is determined during installation and as part of starting SPL.</p>
LOG_MAX_SIZE	<p>Displays the maximum size of the job log file.</p> <p>Once the maximum size is reached, the log file is zipped, and the logs are written into the new file of that job.</p>
RETAIN_LOGS_OF_LAST_DAYS	<p>Displays the number of days up to which the logs are retained.</p>
ENABLE_CERTIFICATE_VALIDATION	<p>Displays true when CA certificate validation is enabled for the host.</p> <p>You can enable or disable this parameter either by editing the <code>spl.properties</code> or by using the SnapCenter GUI or cmdlet.</p>

If any of these parameters are not assigned to the default value or if you want to assign or change the value, then you can modify the `spl.properties` file. You can also verify the `spl.properties` file and edit the file to troubleshoot any issues related to the values that are assigned to the parameters. After you modify the `spl.properties` file, you should restart the SnapCenter Plug-in Loader service.

Steps

1. Perform one of the following actions, as required:

- Start the SnapCenter Plug-in Loader service:
 - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl start`
 - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl start`

- Stop the SnapCenter Plug-in Loader service:
 - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl stop`
 - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl stop`



You can use the `-force` option with the `stop` command to stop the SnapCenter Plug-in Loader service forcefully. However, you should use caution before doing so because it also terminates the existing operations.

- Restart the SnapCenter Plug-in Loader service:
 - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl restart`
 - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl restart`
- Find the status of the SnapCenter Plug-in Loader service:
 - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl status`
 - As a non root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl status`
- Find the change in the SnapCenter Plug-in Loader service:
 - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl change`
 - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl change`

Configure CA certificate with SnapCenter Plug-in Loader (SPL) service on Linux host

You should manage the password of SPL keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to SPL trust-store, and configure CA signed key pair to SPL trust-store with SnapCenter Plug-in Loader service to activate the installed digital certificate.



SPL uses the file 'keystore.jks', which is located at '/var/opt/snapcenter/spl/etc' both as its trust-store and key-store.

Manage password for SPL keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve SPL keystore default password from SPL property file.

It is the value corresponding to the key 'SPL_KEYSTORE_PASS'.

2. Change the keystore password:

```
keytool -storepasswd -keystore keystore.jks
```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

```
keytool -keypasswd -alias "<alias_name>" -keystore keystore.jks
```

Update the same for the key SPL_KEYSTORE_PASS in spl.properties file.

4. Restart the service after changing the password.



>Password for SPL keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to SPL trust-store

You should configure the root or intermediate certificates without the private key to SPL trust-store.

Steps

1. Navigate to the folder containing the SPL keystore: /var/opt/snapcenter/spl/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias <AliasNameForCertificateToBeImported>
-file /<CertificatePath> -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to SPL trust-store.



You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to SPL trust-store

You should configure the CA signed key pair to the SPL trust-store.

Steps

1. Navigate to the folder containing the SPL's keystore /var/opt/snapcenter/spl/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add the CA certificate having both private and public key.

```
keytool -importkeystore -srckeystore <CertificatePathToImport>
-srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
```

5. List the added certificates in the keystore.

```
keytool -list -v -keystore keystore.jks
```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.

7. Change the added private key password for CA certificate to the keystore password.

Default SPL keystore password is the value of the key SPL_KEYSTORE_PASS in spl.properties file.

```
keytool -keypasswd -alias "<aliasNameOfAddedCertInKeystore>" -keystore
keystore.jks
```

8. If the alias name in the CA certificate is long and contains space or special characters ("*", ","), change the alias name to a simple name:

```
keytool -changealias -alias "<OriginalAliasName>" -destalias
"<NewAliasName>" -keystore keystore.jks
```

9. Configure the alias name from the keystore located in spl.properties file.

Update this value against the key SPL_CERTIFICATE_ALIAS.

10. Restart the service after configuring the CA signed key pair to SPL trust-store.

Configure certificate revocation list (CRL) for SPL

You should configure the CRL for SPL

About this task

- SPL will look for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SPL is /var/opt/snapcenter/spl/etc/crl.

Steps

1. You can modify and update the default directory in spl.properties file against the key SPL_CRL_PATH.
2. You can place more than one CRL file in this directory.

The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run `Set-SmCertificateSettings` cmdlet.
- You can display the certificate status for the plug-ins using the `Get-SmCertificateSettings`.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running `Get-Help command_name`. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.



When the status is yellow or green, the data protection operations completes successfully.

Install SnapCenter Plug-in for VMware vSphere

If your database or filesystem is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see [Deployment Overview](#).

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see [Create or import SSL certificate](#).

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is `/opt/netapp/config/crl`.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Prepare for protecting Unix file systems

Before performing any data protection operation such as backup, clone, or restore operations, you should set up your environment. You can also set up the SnapCenter Server to use SnapMirror and SnapVault technology.

To take advantage of SnapVault and SnapMirror technology, you must configure and initialize a data protection relationship between the source and destination volumes on the storage device. You can use NetAppSystem Manager or you can use the storage console command line to perform these tasks.

Before you use the Plug-in for Unix file systems, the SnapCenter administrator should install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server. [Learn more](#)
- Configure the SnapCenter environment by adding storage system connections. [Learn more](#)



SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM registered with SnapCenter using either SVM registration or cluster registration must be unique.

- Add hosts, install the plug-ins, and discover the resources.
- If you are using SnapCenter Server to protect Unix file systems that reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.
- Install Java on your Linux host.
- Configure SnapMirror and SnapVault on ONTAP, if you want backup replication.

Back up Unix file systems

Discover the UNIX file systems available for backup

After installing the plug-in, all the file systems on that host are automatically discovered and displayed in the Resources page. You can add these file systems to resource groups to perform data protection operations.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, and creating storage system connections.
- If the file systems reside on a Virtual Machine Disk (VMDK) or raw device mapping (RDM), you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.

For more information, see [Deploy SnapCenter Plug-in for VMware vSphere](#).

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Path** from the View list.
3. Click **Refresh Resources**.

The file systems are displayed along with information such as type, host name, associated resource groups and policies, and status.

Create backup policies for Unix file systems

Before you use SnapCenter to back up Unix file systems, you must create a backup policy for the resource or the resource group that you want to back up. A backup policy is a set of rules that governs how you manage, schedule, and retain backups. You can also specify the replication, script, and backup type settings. Creating a policy saves time when you want to reuse the policy on another resource or resource group.

Before you begin

- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, discovering the file systems, and creating storage system connections.
- If you are replicating Snapshots to a mirror or vault secondary storage, the SnapCenter administrator must have assigned the SVMs to you for both the source and destination volumes.
- Review the SnapMirror active sync specific prerequisites and limitations. For information, refer [Object limits for SnapMirror active sync](#).

About this task

- SnapLock
 - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.

Specifying a Snapshot locking period prevents deletion of the Snapshots until the retention period expires. This could lead to retaining a larger number of Snapshots than the count specified in the policy.

For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Select **Unix File Systems** from the drop-down list.
4. Click **New**.
5. In the Name page, enter the policy name and details.
6. In the Backup and Replication page, perform the following actions:

- Specify the backup settings.
- Specify the schedule frequency by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.
- In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field...	Do this...
Update SnapMirror after creating a local Snapshot copy	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication). This option should be enabled for SnapMirror active sync.
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups).
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

- In the Retention page, specify the retention settings for the backup type and the schedule type selected in the Backup and Replication page:

If you want to...	Then...
Keep a certain number of Snapshots	<p>Select Copies to keep, and then specify the number of Snapshots that you want to keep.</p> <p>If the number of Snapshots exceeds the specified number, the Snapshots are deleted with the oldest copies deleted first.</p> <p> The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports.</p> <p> You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target.</p>
Keep the Snapshots for a certain number of days	Select Retain copies for , and then specify the number of days for which you want to keep the Snapshots before deleting them.

Snapshot copy locking period	Select Snapshot copy locking period and specify the duration in days, months, or years. Snaplock retention period should be less than 100 years.
------------------------------	--

8. Select policy label.



You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

9. In the Script page, enter the path and the arguments of the prescript or postscript that you want to run before or after the backup operation, respectively.



You should check if the commands exist in the command list available on the plug-in host from the `/_opt/NetApp/snapcenter/scc/etc/allowed_commands.config` path.

You can also specify the script timeout value. The default value is 60 seconds.

10. Review the summary, and then click **Finish**.

Create resource groups and attach policies for Unix file systems


A resource group is a container where you add resources that you want to back up and protect. A resource group allows you to back up all the data that is associated with the file systems.

About this task

- A database with files in ASM disk groups must be in "MOUNT" or "OPEN" state to verify its backups using the Oracle DBVERIFY utility.

Attach one or more policies to the resource group to define the type of data protection job you want to perform.

The following image illustrates the relationship between resources, resource groups, and policies for databases:



- For SnapLock enabled policies, for ONTAP 9.12.1 and below version, if you specify a Snapshot locking

period, the clones created from the tamper proof Snapshots as part of restore will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

- Adding new file systems without SnapMirror active sync to an existing resource group which contains resources with SnapMirror active sync, is not supported.
- Adding new file systems to an existing resource group in failover mode of SnapMirror active sync is not supported. You can add resources to the resource group only in regular or fail-back state.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.

2. In the Resources page, click **New Resource Group**.

3. In the Name page, perform the following actions:

a. Enter a name for the resource group in the Name field.



The resource group name should not exceed 250 characters.

b. Enter one or more labels in the Tag field to help you search for the resource group later.

For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

c. Select the check box, and enter a custom name format that you want to use for the Snapshot name.

For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

4. In the Resources page, select an Unix file systems host name from the **Host** drop-down list.



The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the resources from the Available Resources section and move them to the Selected Resources section.

6. In the Application Settings page, perform the following:

- Select the Scripts arrow and enter the pre and post commands for quiesce, Snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.
- Select one of the backup consistency options:

- Select **File System Consistent** if you want to ensure that file systems cached data is flushed before creating the backup and no input or output operations are allowed on filesystem while creating the backup.



For File System Consistent, Consistency group snapshots will be taken for LUNs involved in Volume group.

- Select **Crash Consistent** if you want to ensure that file systems cached data is flushed before creating the backup.



If you have added different file systems in the resource group, then all volumes from different file systems in the resource group will be put in a Consistency group.

7. In the Policies page, perform the following steps:

- Select one or more policies from the drop-down list.



You can also create a policy by clicking .

In the Configure schedules for selected policies section, the selected policies are listed.

- Click  in the Configure Schedules column for the policy for which you want to configure a schedule.
- In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

Where, *policy_name* is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtpServer`.

9. Review the summary, and then click **Finish**.

Create resource groups and enable secondary protection for Unix file systems on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.

- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

- a. Enter a name for the resource group in the Name field.



The resource group name should not exceed 250 characters.

- b. Enter one or more labels in the Tag field to help you search for the resource group later.

For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

- c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

- d. Specify the destinations of the archive log files that you do not want to back up.



You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.



The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.

6. In the Application Settings page, select the backup option.

7. In the Policies page, perform the following steps:

- a. Select one or more policies from the drop-down list.



You can also create a policy by clicking .

In the Configure schedules for selected policies section, the selected policies are listed.

- b. Click  in the Configure Schedules column for the policy for which you want to configure a schedule.
- c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

Where, *policy_name* is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:

- Select the type of the replication policy.



Synchronous replication policy is not supported.

- Specify the consistency group suffix that you want to use.

- From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.



The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.



If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:

- Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
- Click in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
- In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to...	Do this...
Run verification after backup	Select Run verification after backup .
Schedule a verification	Select Run scheduled verification and then select the schedule type from the drop-down list.

- Select **Verify on secondary location** to verify your backups on secondary storage system.

- Click **OK**.

The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Back up Unix file systems

If a resource is not part of any resource group, you can back up the resource from the Resources page.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, select **Path** from the View list.
3. Click , and then select the host name and the Unix File Systems to filter the resources.
4. Select the file system that you want to back up.
5. In the Resources page, you can perform the following steps:
 - a. Select the check box, and enter a custom name format that you want to use for the Snapshot name.

For example, `customtext_policy_hostname` or `resource_hostname`. A timestamp is appended to the Snapshot name by default.
6. In the Application Settings page, perform the following:
 - Select the Scripts arrow and enter the pre and post commands for quiesce, Snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.
 - Select one of the backup consistency options:
 - Select **File System Consistent** if you want to ensure that file systems cached data is flushed before creating the backup and no operations are performed on filesystem while creating the backup.
 - Select **Crash Consistent** if you want to ensure that file systems cached data is flushed before creating the backup.
7. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.



You can create a policy by clicking .

In the Configure schedules for selected policies section, the selected policies are listed.

- b. Click  in the Configure Schedules column to configure a schedule for the policy you want.
- c. In the Add schedules for policy *policy_name* window, configure the schedule, and then select **OK**.

policy_name is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

8. In the Notification page, select the scenarios in which you want to send the emails from the **Email preference** drop-down list.

You must specify the sender and receiver email addresses, and the subject of the email. If you want to

attach the report of the backup operation performed on the resource, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtpServer`.

9. Review the summary, and then click **Finish**.

The topology page is displayed.

10. Click **Back up Now**.

11. In the Backup page, perform the following steps:

- a. If you have applied multiple policies to the resource, from the Policy drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

- b. Click **Backup**.

12. Monitor the operation progress by clicking **Monitor > Jobs**.

Back up Unix file systems resource groups

You can back up the Unix file systems defined in the resource group. You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups are created according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.

2. In the Resources page, select **Resource Group** from the **View** list.

3. Enter the resource group name in the search box, or click , and select the tag.

Click  to close the filter pane.

4. In the Resource Group page, select the resource group to back up.

5. In the Backup page, perform the following steps:

- a. If you have multiple policies associated with the resource group, select the backup policy you want to use from the **Policy** drop-down list.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

- b. Select **Backup**.

6. Monitor the progress by selecting **Monitor > Jobs**.

Monitor Unix file systems backup

Learn how to monitor the progress of backup operations and data protection operations.

Monitor Unix file systems backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
 - a. Click  to filter the list so that only backup operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Backup**.
 - d. From the **Status** drop-down, select the backup status.
 - e. Click **Apply** to view the operations completed successfully.
4. Select a backup job, and then click **Details** to view the job details.



Though the backup job status displays  , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click  on the Activity pane to view the five most recent operations.

When you click one of the operations, the operation details are listed in the **Job Details** page.

View protected Unix file systems in the Topology page

When you are preparing to back up, restore, or clone a resource, you might find it helpful to view a graphical representation of all backups, restored file systems, and clones on the primary and secondary storage.

About this task

In the Topology page, you can see all the backups, restored file systems, and clones that are available for the selected resource or resource group. You can view the details of those backups, restored file systems, and clones, and then select them to perform data protection operations.

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

-  displays the number of backups and clones that are available on the primary storage.
-  displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
-  displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.



Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

If you have secondary relationship as SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]), you can see following additional icons:

-  The replica site is up.
-  The replica site is down.
-  The secondary mirror or vault relationship has not been re-established.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.

3. Select the resource either from the resource details view or from the resource group details view.

If the resource is protected, the Topology page of the selected resource is displayed.

4. Review the Summary card to see a summary of the number of backups and clones available on the primary and secondary storage.

The Summary Card section displays the total number of backups and clones.

Clicking the **Refresh** button starts a query of the storage to display an accurate count.

If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

When the file system is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

For SnapMirror active sync, clicking the **Refresh** button refreshes the SnapCenter backup inventory by querying ONTAP for both primary and replica sites. A weekly schedule also performs this activity for all databases containing SnapMirror active sync relationship.

- For SnapMirror active sync and only for ONTAP 9.14.1, Async Mirror or Async MirrorVault relationships to the new primary destination should be manually configured after failover. From ONTAP 9.15.1 onwards Async Mirror or Async MirrorVault is auto configured to the new primary destination.
- After failover, a backup should be created for SnapCenter to be aware of the failover. You can click **Refresh** only after a backup has been created.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, and delete operations.



You cannot rename or delete backups that are on the secondary storage.

7. If you want to delete a clone, select the clone from the table, and then click .

Example showing backups and clones on the primary storage



Restore and recover Unix file systems

Restore Unix file systems

In the event of data loss, you can use SnapCenter to restore Unix file systems.

About this task

- You should execute the following commands to establish the connection with the SnapCenter Server, list the backups and retrieve its information, and restore the backup.

The information regarding the parameters that can be used with the command and their descriptions can be obtained by running `Get-Help command_name`. Alternatively, you can also refer to the [SnapCenter Software Command Reference Guide](#).

- For SnapMirror active sync restore operation, you must select the backup from the primary location.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Path** or **Resource Group** from the **View** list.
3. Select the file system from either the details view or the resource group details view.

The topology page is displayed.

4. From the Manage Copies view, select **Backups** from either the primary or the secondary (mirrored or replicated) storage systems.
5. Select the backup from the table, and then click .
6. In the Restore Scope page:
 - For NFS file systems, by default **Connect and Copy** restore is selected. You can also select **Volume Revert** or **Fast Restore**.
 - For non NFS file systems, the restore scope is selected depending on the layout.
7. In the PreOps page, enter pre restore commands to run before performing a restore job.
8. In the PostOps page, enter post restore commands to run after performing a restore job.



You should check if the commands exist in the command list available on the plug-in host at location `/opt/NetApp/snapcenter/scc/etc/allowed_commands.config` path.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the email notifications.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the restore operation performed, you must select **Attach Job Report**.



For email notification, you must have specified the SMTP server details by using the either the GUI or the PowerShell command `Set-SmSmtpServer`.

10. Review the summary, and then click **Finish**.



If restore operation fails, rollback is not supported.



In case of restore of a filesystem residing on volume group, the old contents on the filesystem are not deleted. Only the content from the cloned filesystem will be copied to the source filesystem. This is applicable when there are multiple filesystems on the volume group and default NFS filesystem restores.

11. Monitor the operation progress by clicking **Monitor > Jobs**.

Monitor Unix file systems restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
 - a. Click  to filter the list so that only restore operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Restore**.
 - d. From the **Status** drop-down list, select the restore status.
 - e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

The **View logs** button displays the detailed logs for the selected operation.

Clone Unix file systems

Clone Unix file system backup

You can use SnapCenter to clone Unix file system using the backup of the filesystem.

Before you begin

- You can skip the fstab file update by setting the value of `SKIP_FSTAB_UPDATE` to `true` in the `agent.properties` file located at `/opt/NetApp/snapcenter/scc/etc`.
- You can have a static clone volume name and junction path by setting the value of `USE_CUSTOM_CLONE_VOLUME_NAME_FORMAT` to `true` in the `agent.properties` file located at `/opt/NetApp/snapcenter/scc/etc`. After updating the file, you should restart the SnapCenter Plug-in creator service by running the command: `/opt/NetApp/snapcenter/scc/bin/scc restart`.

Example: Without this property the clone volume name and junction path will be like

`<Source_volume_name>_Clone_<Timestamp>` but now it will be
`<Source_volume_name>_Clone_<Clone_Name>`

This keeps the name constant so that you can manually keep the fstab file updated if you do not prefer to update the fstab by SnapCenter.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Path** or **Resource Group** from the **View** list.
3. Select the file system from either the details view or the resource group details view.

The topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field...	Do this...
Clone server	By default, the source host is populated.
Clone mount point	Specify the path where the file system will be mounted.

7. In the Scripts page, perform the following steps:

- a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.



You should check if the commands exist in the command list available on the plug-in host from the `/opt/NetApp/snapcenter/scc/etc/allowed_commands.config` path.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the clone operation performed, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtpServer`.

9. Review the summary, and then click **Finish**.
10. Monitor the operation progress by clicking **Monitor > Jobs**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, [Split a FlexClone volume from its parent volume](#).
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select Database from the View list.
For file systems	Select Path from the View list.

3. Select the appropriate resource from the list.

The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor > Jobs**.

The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

For example:

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```

The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

[SnapCenter clone or verification fails with aggregate does not exist](#)

Monitor Unix file systems clone operations

You can monitor the progress of SnapCenter clone operations by using the **Jobs** page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the **Jobs** page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
 - a. Click to filter the list so that only clone operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Clone**.
 - d. From the **Status** drop-down list, select the clone status.
 - e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Copyright information

Copyright © 2026 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.