



SnapCenter® Software 4.1

Data Protection Guide

For Oracle® Databases

August 2019 | 215-13397_2019-08_en-us
doccomments@netapp.com

Updated for 4.1.1



Contents

Deciding whether to read the SnapCenter Data Protection Guide for Oracle Database	6
SnapCenter Plug-in for Oracle Database overview	7
Data protection workflow for Oracle databases	8
Preparing for data protection	9
Prerequisites for using the SnapCenter Plug-in for Oracle Database	9
Storage types supported by SnapCenter Plug-in for Oracle Database	10
How resources, resource groups, and policies are used for protecting Oracle databases	11
Logging in to SnapCenter	12
Backing up Oracle databases	14
Configuring Run As account credentials for an Oracle database	15
Determining whether Oracle databases are available for backup	16
Creating backup policies for Oracle databases	17
Creating resource groups and attaching policies for Oracle databases	21
Backing up Oracle resources	23
Backing up Oracle database resource groups	25
Monitoring backup operations	26
Monitoring operations in the Activity pane	27
Canceling the SnapCenter Plug-in for Oracle Database backup operations	27
Viewing Oracle database backups and clones in the Topology page	28
Mounting and unmounting database backups	30
Mounting a database backup	30
Unmounting a database backup	31
Restoring Oracle databases	32
Restoring an Oracle database	32
Monitoring restore operations	35
Canceling restore operations	36
Cloning Oracle database	38
Cloning an Oracle database	39
Monitoring clone operations in SnapCenter	45
Canceling clone operations	46
Splitting an Oracle Database Clone	46
Backing up, restoring, and cloning using Linux commands	48
Backing up Oracle databases using Linux commands	48
Restoring and recovering Oracle databases using Linux commands	49
Cloning Oracle database backups using Linux commands	50
Refreshing a clone	50
Troubleshooting data protection operations	52
Discovery operation takes long time to complete	52
Unable to add Linux host to SnapCenter	52

Scanning of host bus adapters takes long time to complete	52
Backup fails during the discovery of file system on a VM	53
Backup operation fails during the storage discovery process	53
Backup operation fails if database query is timed out	53
Backup operation might fail when the OS group ID of the Oracle database administrator is changed	54
Backup fails with error: Storage system(s) may need to be added, also ensure that the associated host is in a connected state	55
Backup operation might fail if external RMAN catalog database has issues	55
Cataloging and uncataloging with Oracle RMAN will fail if the execution time is beyond the timeout value	56
Unable to find Snapshot copy after successfully creating the backup	56
Backup operation fails if Snapshot copies on the secondary storage reaches maximum limit	57
Unable to update the SnapMirror and SnapVault status	57
ASM backup verification fails	57
Backup verification fails when files are not accessible	58
Backup verification is timed out	58
Disk paths are not included in the asm_diskstring database parameter	59
Failed to mount ASM log backups as part of recovery operation	59
Unable to change the database state from shutdown to mount	59
Restore operation of datafiles and control files fail	60
Restore from a secondary SnapMirror or SnapVault location fails	60
Restore operation might fail if the database size is in terabytes	61
Restore operation fails when you select a backup of an orphan incarnation	61
Clone operation will fail if multipath is disabled on the plug-in host	62
Cloning operation will fail in SAN environments on OL 7 or later or RHEL 7 or later	62
Clone operation will fail due to inaccessible virtual device	63
Clone operation might fail or take longer time to complete with default TCP_TIMEOUT value	63
Clone operation might fail if you are using Oracle databases 11.2.0.3 or later	64
Fails to clone an Oracle database on a volume	64
Recovery of a cloned database fails	64
Recovery operation fails if the SCN specified is inconsistent	65
Clone split operation stops responding	65
Clone split estimation fails when the aggregate does not have space	65
Clone split start operation fails	65
Databases on which the clone split operation was performed are listed as clones ...	66
File system is not deleted during the clone delete operation	66
Backup and clone operations fail if stale entries of the cloned disk group exists	67
Operations fail when there is insufficient space to create Snapshot copies	67
Operations are not executed due to insufficient space in the root file system	67
Data protection operation fails if operational lock file is not deleted	68
Data protection operation fails because of application firewall	68

Operations that require backup to be mounted might fail	69
Messages in the log file display incorrect time zone	69
Operations fail with command execution timeout error	69
Data protection operation fails in a non-multipath environment in RHEL 7 and later	70
Managing policies	71
Detaching policies	71
Modifying policies	72
Deleting policies	73
Managing resource groups	74
Stopping and resuming operations on resource groups	74
Deleting resource groups	75
Managing backups	76
Renaming backups	76
Deleting backups	76
Managing clones	78
Deleting clones	78
Copyright	79
Trademark	80
How to send comments about documentation and receive update notifications	81

Deciding whether to read the SnapCenter Data Protection Guide for Oracle Database

This guide describes how to use SnapCenter to perform backup, restore, and clone operations on Oracle databases.

You should read this information if you want to use SnapCenter in the following ways:

- You want to create data protection policies and resource groups for Oracle databases
- You want to perform backup, restore, or clone operations on Oracle databases using the graphical user interface (GUI)
- You want to perform backup, restore, or clone operations on Oracle databases using Linux-based commands

You should have already performed the following tasks:

- Installed SnapCenter Server and the SnapCenter Plug-ins Package for Linux
- Configured role-based access control (RBAC), storage system connections, and Run As accounts
- Installed the SnapCenter Plug-in for VMware vSphere and registered it with SnapCenter, if you want to protect Oracle databases on virtual machines
- Set up SnapMirror and SnapVault relationships, if you want backup replication

You can also use the following information to help accomplish your data protection goals:

- SnapCenter Server and plug-in installation and setup
[*Installing and setting up SnapCenter*](#)
[*Getting Started*](#)
- SnapCenter concepts, including architecture, features, and benefits
[*Concepts*](#)
- Other SnapCenter Data Protection Guides for specific types of resources, such as Microsoft SQL Server, Oracle, Windows file systems, and custom plug-ins
- SnapCenter Linux commands
[*SnapCenter Software 4.1 Linux Command Reference Guide*](#)
- SnapCenter administration, including dashboards, reporting capabilities, and REST APIs, and managing licenses, storage connections, and the SnapCenter Server repository
[*Performing administrative tasks*](#)

SnapCenter Plug-in for Oracle Database overview

The SnapCenter Plug-in for Oracle Database is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Oracle databases. The Plug-in for Oracle Database automates the backup, cataloging and uncataloging with Oracle RMAN, verification, mounting, unmounting, restore, recovery, and cloning of Oracle databases in your SnapCenter environment.

The Plug-in for Oracle Database installs SnapCenter Plug-in for UNIX to perform all the data protection operations.

You can use the Plug-in for Oracle Database to manage backups of Oracle databases running SAP applications. However, SAP BR*Tools integration is not supported.

For information about the SnapCenter architecture, features, and benefits, see the SnapCenter concepts documentation.

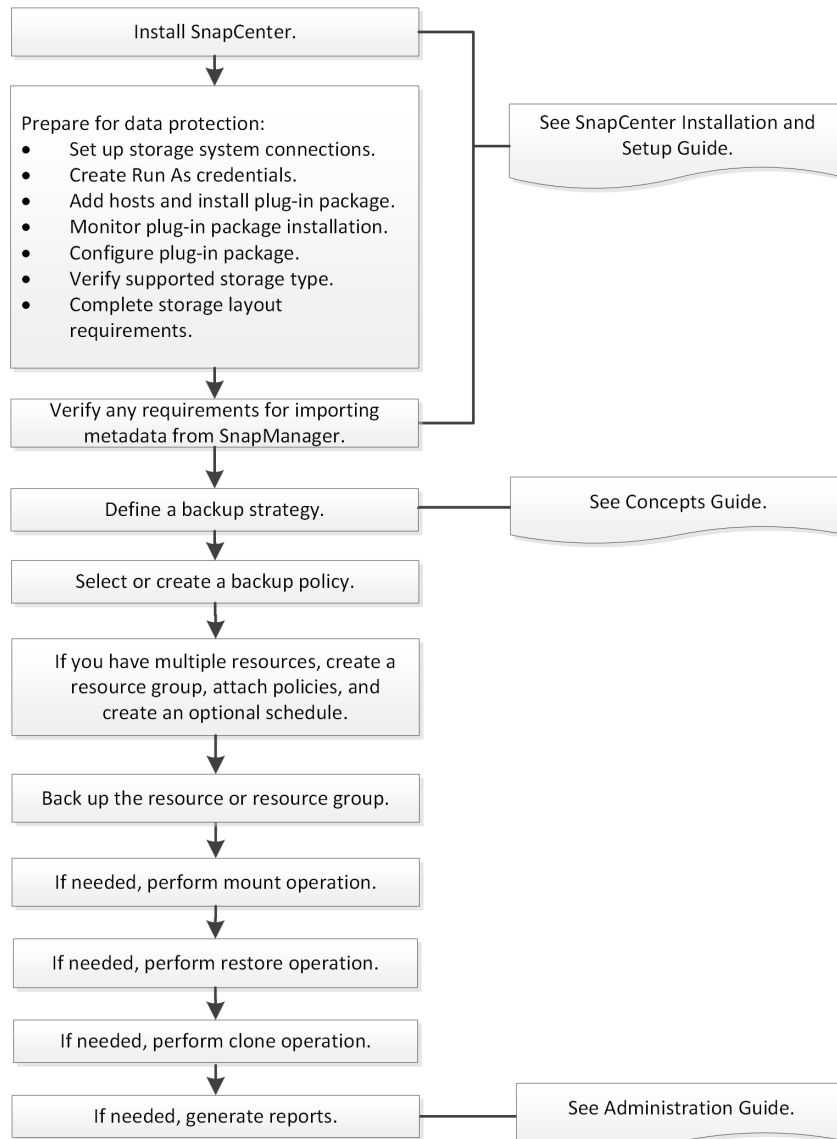
Related information

[*Concepts*](#)

[*Installing and setting up SnapCenter*](#)

Data protection workflow for Oracle databases

The data protection workflow lists the tasks that you have to perform for data protection.



Related tasks

[Backing up Oracle databases](#) on page 14

[Mounting and unmounting database backups](#) on page 30

[Restoring Oracle databases](#) on page 32

[Cloning an Oracle database](#) on page 39

Related information

[Installing and setting up SnapCenter](#)

[Concepts](#)

[Performing administrative tasks](#)

Preparing for data protection

Before performing any data protection operation such as backup, clone, or restore operations, you must define your strategy and set up the 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 NetApp ONTAP System Manager or you can use the storage console command line to perform these tasks.

Prerequisites for using the SnapCenter Plug-in for Oracle Database

Before you use the Plug-in for Oracle Database, the SnapCenter administrator must install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter Server.
- Configure the SnapCenter environment by adding storage system connections.
 - Note:** SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM supported by SnapCenter must have a unique name.
- Create Run As with authentication mode as Linux for the install user.
- Add hosts, install the plug-ins, and discover the resources.
- If you are using SnapCenter Server to protect Oracle databases that reside on VMware RDM LUNs or VMDKs, you must install the SnapCenter Plug-in for VMware vSphere.
- Install Java on your Linux host.

For the latest information about supported versions, see the NetApp Interoperability Matrix Tool (IMT).
- If you have Oracle databases on NFS environments, you must have configured at least one NFS data LIF for primary or secondary storage to perform mount, clone, verification, and restore operations.
- If you have multiple data paths (LIFs) or a dNFS configuration, you can perform the following using the SnapCenter CLI on the database host:
 - By default, all the IP addresses of the database host are added to the NFS storage export policy in storage virtual machine (SVM) for the cloned volumes. If you want to have a specific IP address or restrict to a subset of the IP addresses, run the `Set-PreferredHostIPsInStorageExportPolicy` CLI.
 - If you have multiple data paths (LIFs) in SVM, SnapCenter chooses the appropriate data path (LIF) for mounting the NFS cloned volume. However, if you want to specify a specific data path (LIF), you must run the `Set-SvmPreferredDataPath` CLI.

For more information, see the Linux command reference information.

- If you have Oracle databases on SAN environments, ensure that the SAN environment is configured as per the recommendation mentioned in *Recommended Host Settings for Linux Unified Host Utilities* and *Using Linux Hosts with ONTAP storage* guides.

- If you are have Oracle databases on LVM in Oracle Linux or RHEL operating systems, install the latest version of Logical Volume Management (LVM).
For the latest information about supported versions, see the NetApp Interoperability Matrix Tool (IMT).
- If you are using SnapManager for Oracle and want to SnapCenter Plug-in for Oracle Database, you can migrate the profiles to policies and resource groups of SnapCenter by using the `sccli` command `sc-migrate`.
- Configure SnapMirror and SnapVault on ONTAP, if you want backup replication

Related information

[NetApp Interoperability Matrix Tool](#)

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

[Recommended Host Settings for Linux Unified Host Utilities 7.1](#)

[Using Linux Hosts with ONTAP Storage](#)

[Installing and setting up SnapCenter](#)

Storage types supported by SnapCenter Plug-in for Oracle Database

SnapCenter supports a wide range of storage types on both physical and virtual machines. You must verify the support for your storage type before installing the SnapCenter Plug-ins Package for Linux.

Storage provisioning using SnapCenter is not supported for SnapCenter Plug-ins Package for Linux.

Machine	Storage type
Physical server	FC-connected LUNs
	iSCSI-connected LUNs
	NFS-connected volumes
VMware ESX	RDM LUNs connected by an FC or iSCSI ESXi HBA
	iSCSI LUNs connected directly to the guest system by the iSCSI initiator
	VMDKs on VMFS or NFS datastores
	NFS volumes connected directly to the guest system

The storage configurations not supported by SnapCenter Plug-in for Oracle Database are:

- ASM on VMDK
- ASM on RDM LUN
- RAC on VMDK
- RAC on RDM LUN

How resources, resource groups, and policies are used for protecting Oracle databases


Before you use SnapCenter, it is helpful to understand basic concepts related to the backup, clone, and restore operations you want to perform.

- *Resources* are typically Oracle databases in SnapCenter.
- A SnapCenter *resource group* is a collection of resources on a host or cluster.
When you perform an operation on a resource group, you perform that operation on the *resources* defined in the resource group.
The resource groups were formerly known as *datasets*.
- The *policies* specify the backup frequency, copy retention, replication, scripts, and other characteristics of data protection operations.
When you create a resource group, you select one or more policies for that group. You can also select a policy when you perform a backup on demand for a single resource. You can also perform scheduled backups for single resources and resource groups.

Think of a resource group as defining *what* you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining *how* you want to protect it.

If you are backing up all databases of a host, for example, you might create a resource group that includes all of the databases in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a full backup daily and another schedule that performs log backups hourly.

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



Logging in to SnapCenter

SnapCenter supports role-based access control (RBAC). SnapCenter administrator assigns a role and resources through SnapCenter RBAC for user in workgroup or active directory and groups in active directory. The RBAC user can then log in to SnapCenter with the assigned roles.

Before you begin

- You should have enabled Windows Process Activation Service (WAS) in Windows Server Manager.
- If you want to use Internet Explorer as the browser to log in to the SnapCenter Server, you the Protected Mode in Internet Explorer must be disabled.

About this task

During the installation, the SnapCenter Server Install wizard creates a shortcut and places it on the desktop and in the Start menu of the host where SnapCenter is installed. Additionally, at the end of the installation, the Install wizard displays the SnapCenter URL based on information you supplied during the installation, which you can copy if you want to log in from a remote system.

Attention: Closing just the SnapCenter browser tab does not log you out of SnapCenter if you have multiple tabs open in your web browser. To end your connection with SnapCenter, you must log out of SnapCenter either by clicking the **Sign out** button or closing the entire web browser.

Best Practice: For security reasons, it is recommended that you not allow your browser to save your SnapCenter password.

The default GUI URL is a secure connection to the default port 8146 on the server where the SnapCenter Server is installed (`https://server:8146`). If you provided a different server port during the SnapCenter installation, that port is used instead.

For Network Load Balance (NLB) deployment, you must access SnapCenter using the NLB cluster IP (`https://NLB_Cluster_IP:8146`). If you do not see the SnapCenter UI when you navigate to `https://NLB_Cluster_IP:8146` in Internet Explorer (IE), you must add the NLB IP address as a trusted site in IE on each plug-in host, or you must disable IE Enhanced Security on each plug-in host.

[NetApp KB Article 2025082: SnapCenter in an HA configuration with Application Request Routing enabled.](#)

In addition to using the SnapCenter GUI, you can use PowerShell cmdlets to script configuration, backup, and restore operations.

Note: Some cmdlets might have changed. If you use cmdlets in older versions of SnapCenter scripts, you might need to update your scripts.

The SnapCenter cmdlet or SnapCenter CLI documentation has the details.

Note: If you are logging in to SnapCenter for the first time, you must log in using the credentials that you provided during the install process.

Steps

1. Launch SnapCenter from the shortcut located on your local host desktop, from the URL provided at the end of the installation, or from the URL provided to you by your SnapCenter administrator.
2. Enter your user credentials.

To specify the following...	Use one of these formats...
Domain administrator	<i>NetBIOS\UserName</i> <i>UserName@UPN suffix</i> For example, username@netapp.com <i>Domain FQDN\UserName</i>
Local administrator	<i>UserName</i>

3. If you are assigned more than one role, from the **Role** box, select the role that you want to use for this login session.

Your current user and associated role are shown in the upper right of SnapCenter after you are logged in.

Result

If you are using SnapCenter for the first time, the Storage Systems page is displayed, and the Get Started pane is expanded.

If the login fails with the error that site cannot be reached, you should map the SSL certificate to SnapCenter.

After you finish

If you have untrusted Active Directory domains that you want SnapCenter to support, you must register those domains with SnapCenter before configuring the roles for the users on untrusted domains.

[Performing administrative tasks](#)

Related information

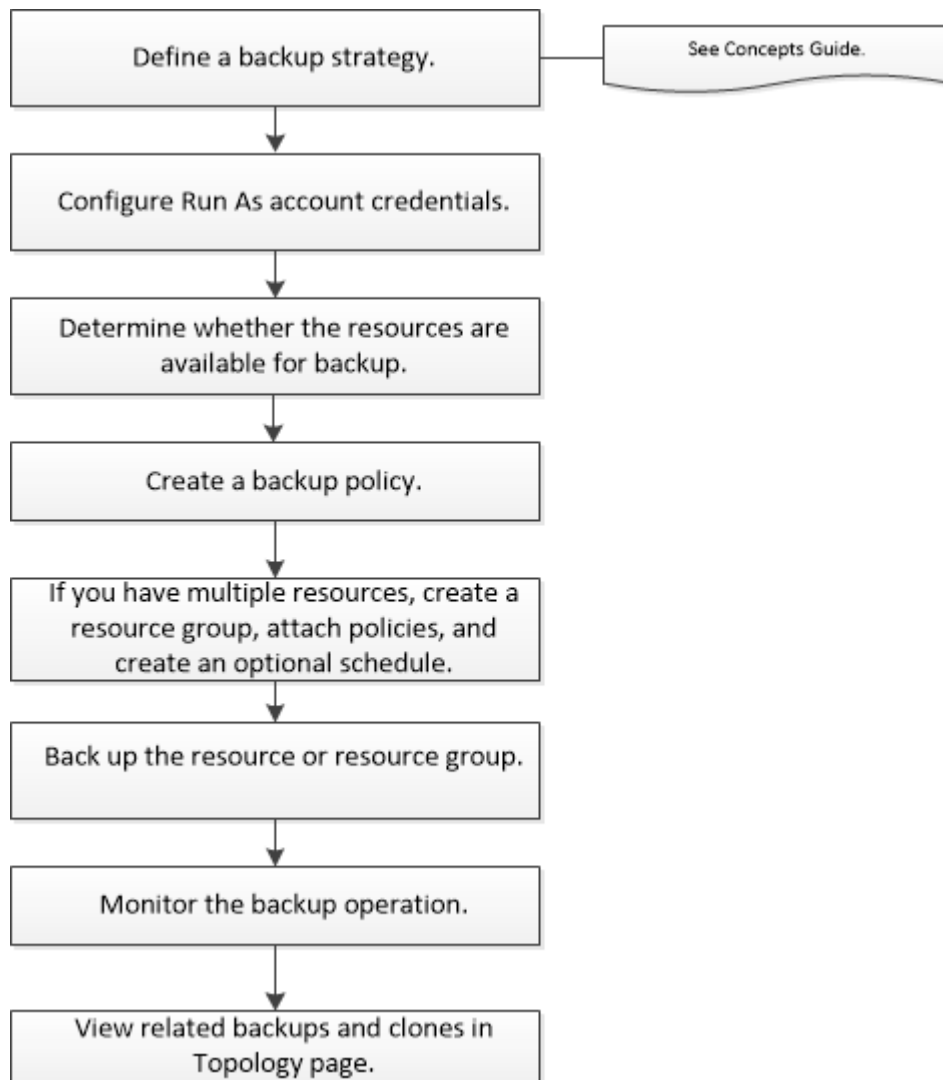
[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Backing up Oracle databases

You can either create a backup of a resource (database) or resource group. The backup workflow includes planning, identifying the resources for backup, creating backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

About this task

The following workflow shows the sequence in which you must perform the backup operation:



Note: While creating a backup for Oracle databases, an operational lock file (`.sm_lock_dbid`) is created on the Oracle database host in the `$ORACLE_HOME/dbs` directory to avoid multiple operations being executed on the database. After the database has been backed up, the operational lock file is automatically removed.

You can also use Linux commands manually or in scripts to perform backup operations. For detailed information about Linux commands, use the SnapCenter command help or see the command reference information.

Related tasks

[Backing up Oracle databases using Linux commands](#) on page 48

Related information

[Concepts](#)

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Configuring Run As account credentials for an Oracle database

You must configure Run As account credentials that are used to perform data protection operations on Oracle databases.

Before you begin



If you set up Run As credentials for individual resource groups and the user name does not have full admin privileges, the user name must at least have resource group and backup privileges.

About this task

If you have enabled an Oracle database authentication, a red lock icon is shown in the resources view. You must configure database Run As credentials to be able to protect the database or add it to the resource group to perform data protection operations.

Note: If you specify incorrect credentials while creating a Run As account, an error message is displayed. You must click **Cancel**, and then retry.


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 **Database** from the **View** list.
Click , and select the host name and the database type to filter the resources. You can then click  to close the filter pane.
3. Select the database, and then click **Database Settings > Configure Database**.
4. In the **Configure database settings** section, from the **Use existing Run As** drop-down list, select the Run As account that should be used to perform data protection jobs on the Oracle database.

You can also create a Run As credential by clicking .

Note: The database port number is automatically populated.

5. In the **Configure ASM settings** section, from the **Use existing Run As** drop-down list, select the Run As account that should be used to perform data protection jobs on the ASM instance.

You can also create a Run As credential by clicking .

Note: The ASM instance port number is automatically populated.

6. In the **Configure RMAN catalog settings** section, from the **Use existing Run As** drop-down list, select the Run As account that should be used to perform data protection jobs on the Oracle Recovery Manager (RMAN) catalog database.

You can also create a Run As credential by clicking .

In the **TNSName** field, enter the Transparent Network Substrate (TNS) file name that will be used by the SnapCenter Server to communicate with the database.

7. In the **Preferred RAC Nodes** field, specify the Real Application Cluster (RAC) nodes preferred for backup.

The preferred nodes might be one or all cluster nodes where the RAC database instances are present. The backup operation is triggered only on these preferred nodes in the order of preference.

In RAC One Node, only one node is listed in the preferred nodes, and this preferred node is the node where the database is currently hosted.

After failover or relocation of RAC One Node database, refreshing of resources in the SnapCenter Resources page will remove the host from the **Preferred RAC Nodes** list where the database was earlier hosted. The RAC node where the database is relocated will be listed in **RAC Nodes** and will need to be manually configured as the preferred RAC node.

8. Click **OK**.

Determining whether Oracle databases are available for backup

Resources are Oracle databases on the host that are maintained by SnapCenter. You can add these databases to resource groups to perform data protection operations after you discover the databases that are available.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, creating storage system connections, and adding credentials.
- If the databases reside on a Virtual Machine Disk (VMDK) or raw device mapping (RDM), you must have installed SnapCenter Plug-in for VMware vSphere.
- If databases reside on a VMDK file system, you must have logged in to vCenter and navigated to **VM options > Advanced > Edit configuration** to set the value of `disk.enableUUID` to **true** for the VM.
- You must have reviewed the process that SnapCenter follows to discover different types and versions of Oracle databases.

[Concepts](#)



About this task

After installing the plug-in, all of the databases on that host are automatically discovered and displayed in the Resources page.

The databases should be at least in the mounted state or above for the discovery of the databases to be successful. In an Oracle Real Application Clusters (RAC) environment, the RAC database instance in the host where the discovery is performed, should be at least in the mounted state or above for the discovery of the database instance to be successful. Only the databases that are discovered successfully can be added to the resource groups.

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 **Database** from the **View** list.

Click the  icon, and then select the host name and the database type to filter the resources. You can then click the  icon to close the filter pane.

3. Click **Refresh Resources**.

In a RAC One Node scenario, the database is discovered as the RAC database on the node where it is currently hosted.

Result

The databases are displayed along with information such as database type, host or cluster name, associated resource groups and policies, and status.

- If the database is on a non-NetApp storage system, the user interface displays a `Not available for backup` message in the Overall Status column.
You cannot perform data protection operations on the database that is on a non-NetApp storage system.
- If the database is on a NetApp storage system and not protected, the user interface displays a `Not protected` message in the Overall Status column.
- If the database is on a NetApp storage system and protected, the user interface displays an `Available for backup` message in the Overall Status column.

Note: If you have enabled an Oracle database authentication, a red color lock icon is shown in the resources view. You must configure database credentials to be able to protect the database or add it to the resource group to perform data protection operations.

Creating backup policies for Oracle databases

Before you use SnapCenter to back up Oracle database resources, 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 defined your backup strategy.
For details, see the information about defining a data protection strategy for Oracle databases.
[Concepts](#)
- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, discovering databases, and creating storage system connections.
- If you are replicating Snapshot copies to a mirror or vault secondary storage, the SnapCenter administrator must have assigned the SVMs to you for both the source and destination volumes.
For information about how administrators assign resources to users, see the SnapCenter installation information.
[Installing and setting up SnapCenter](#)

Steps

1. In the left navigation pane, click **Settings**.
2. In the **Settings** page, click **Policies**.
3. Select **Oracle Database** from the drop-down list.

4. Click **New**.
5. In the **Name** page, enter the policy name and description.
6. In the **Backup Type** page, perform the following steps:
 - If you want to **create an online backup**, select **Online backup**.
You must specify whether you want to back up all the database files, only datafiles and control files, or only archive log files.

Note: Online backup of a container database (CDB) fails if one of the PDBs of the CDB has been created by cloning an existing PDB, and if the PDB is in mount state.
 - If you want to **create an offline backup**, select **Offline backup**, and then select one of the following options:
 - If you want to create an offline backup when the database is in mounted state, select **Mount**.
 - If you want to create an offline shutdown backup by changing the database to shutdown state, select **Shutdown**.
If you are using Oracle 12c database, and want to save the state of the pluggable databases (PDBs) before creating the backup, you must select **Save state of PDBs**. This enables you to bring the PDBs to their original state after the backup is created.
 - Specify the schedule frequency by selecting **Hourly**, **Daily**, **Weekly**, or **Monthly**.
Note: You can specify the schedule (start date and end date) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but enables you to assign different backup schedules to each policy.
 - If you want to catalog backup using Oracle Recovery Manager (RMAN), select **Catalog backup with Oracle Recovery Manager (RMAN)**.
You can perform deferred cataloging for one backup at a time only by using the SnapCenter CLI command `Catalog-SmBackupWithOracleRMAN`.
Note: If you want to catalog backups of a RAC database, ensure that no other job is running for that database. If another job is running, the cataloging operation fails instead of getting queued.
 - If you want to prune archive logs after backup, select **Prune archive logs after backup**.
Note: Pruning of archive logs from the archive log destination that is unconfigured in the database, will be skipped.

You can delete archive logs only if you have selected the archive log files as part of your backup.

Note: You must ensure that all the nodes in an RAC environment can access all the archive log locations for the delete operation to be successful.

If you want to...	Then...
Delete all archive logs	Select Delete all archive logs .
Delete archive logs that are older	Select Delete archive logs older than , and then specify the age of the archive logs that are to be deleted in days and hours.
Delete archive logs from all destinations	Select Delete archive logs from all the destinations .

If you want to...	Then...
Delete the archive logs from the log destinations that are part of the backup	Select Delete archive logs from the destinations which are part of backup .

☒ Prune archive logs after backup

Prune log retention setting

☐ Delete all archive logs

☒ Delete archive logs older than

Prune log destination setting

☐ Delete archive logs from all the destinations

☒ Delete archive logs from the destinations which are part of backup

7. In the **Retention** page, specify the retention settings for the backup type and the schedule type selected in the **Backup Type** page:

If you want to...	Then...
Keep a certain number of Snapshot copies	<p>Select Total Snapshot copies to keep, and then specify the number of Snapshot copies that you want to keep.</p> <p>If the number of Snapshot copies exceeds the specified number, the Snapshot copies are deleted with the oldest copies deleted first.</p> <p>Note: The maximum retention value is 1018 for resources on ONTAP 9.4 or later, and 254 for resources on ONTAP 9.3 or earlier. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports.</p> <p>Important: 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 copy is the reference Snapshot copy for the SnapVault relationship until a newer Snapshot copy is replicated to the target.</p>
Keep the Snapshot copies for a certain number of days	Select Keep Snapshot copies for , and then specify the number of days for which you want to keep the Snapshot copies before deleting them.

Note: You can retain archive log backups only if you have selected the archive log files as part of your backup.

8. In the **Replication** page, specify the replication settings:

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).
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups).

For this field...	Do this...
Secondary policy label	Select a Snapshot label. Depending on the Snapshot label that you select, the ONTAP software applies the secondary Snapshot copy retention policy that determines how Snapshot copies are retained on the secondary storage system.
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

Select secondary replication options

☐ Update SnapMirror after creating a local Snapshot copy.

☐ Update SnapVault after creating a local Snapshot copy.

Secondary policy label

One Time



Error retry count

3



9. Optional: 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 must store the prescripts and postscripts either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

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

10. In the **Verification** page, perform the following steps:

- a. Select the backup schedule for which you want to perform the verification operation.
- b. In the **Verification script** commands section, enter the path and the arguments of the prescript or postscript that you want to run before or after the verification operation, respectively.

You must store the prescripts and postscripts either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

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

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

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

[Installing and setting up SnapCenter](#)

[Concepts](#)

Creating resource groups and attaching policies for Oracle databases

A resource group is the container to which you must add resources that you want to back up and protect. A resource group enables you to back up all the data that is associated with a given application simultaneously. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

Before you begin

You should ensure that the database having files on the ASM disk groups should be either in “MOUNT” or “OPEN” state to verify its backups using the Oracle DBVERIFY utility.


Steps

1. In the left navigation pane, click **Resources**, and then select 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:


For this field...	Do this...
Name	Enter a name for the resource group.
Tags	Enter one or more labels that will help you later search for the resource group. 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.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot copy name. For example, <i>customtext_resource_group_policy_hostname</i> or <i>resource_group_hostname</i> . By default, a timestamp is appended to the Snapshot copy name.
Exclude archive log destinations from backup	Specify the destinations of the archive log files that you do not want to back up.

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

Note: 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 then click the right arrow to move them to the **Added** section.
6. In the **Policies** page, perform the following steps:
 - a. Select one or more policies from the drop-down list.

Note: 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.

7. In the **Verification page, perform the following steps:**

- a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
- b. Click  in the **Configure Schedules** column to configure the verification schedule for all the schedule types of the policy.
- c. 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.

- d. Select **Verify on secondary location** to verify your backups on secondary storage system.
- e. Click **OK**.

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

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**.

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

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

Related tasks

[Creating backup policies for Oracle databases](#) on page 17

Related information

[Concepts](#)

Backing up Oracle resources



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

Before you begin

- You must have created a resource group with a policy attached.
- If you are backing up to a resource that has a SnapMirror relationship to secondary storage, you must have SnapCenter admin privileges and the scadmin role must include the “snapmirror all” privilege.
- You must have assigned the aggregate that is being used by the backup operation to the storage virtual machine (SVM) used by the database.
- You should ensure that all data volumes and archive log volumes belonging to the database are protected if secondary protection is enabled for that database.
- You should ensure that the database having files on the ASM disk groups should be either in “MOUNT” or “OPEN” state to verify its backups using the Oracle DBVERIFY utility.

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 **Database** from the **View** list.

Click  and select the host name and the database type to filter the resources. You can then click  to close the filter pane.


3. Click the database that you want to back up.

The Database-Protect page is displayed.


4. In the **Resource** page, perform the following actions:

For this field...	Do this...
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot copy name. For example, <i>customtext__policy_hostname</i> or <i>resource_hostname</i> . By default, a timestamp is appended to the Snapshot copy name.
Exclude archive log destinations from backup	Specify the destinations of the archive log files that you do not want to back up.

5. In the **Policies** page, perform the following steps:
 - a. Select one or more policies from the drop-down list.

Note: 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.

6. In the **Verification page, perform the following steps:**

- a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
- b. Click  in the **Configure Schedules** column to configure the verification schedule for all the schedule types of the policy.
- c. 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.

- d. Select **Verify on secondary location** to verify your backups on secondary storage system.
- e. Click **OK**.

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

7. 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 backup operation performed on the resource, select **Attach Job Report**.

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

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

The database topology page is displayed.

9. Click **Back up Now.**

10. 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.
- b. Click **Backup**.

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

Related tasks

[Creating backup policies for Oracle databases](#) on page 17

[Monitoring backup operations](#) on page 26

[Viewing Oracle database backups and clones in the Topology page](#) on page 28

[Backing up Oracle databases using Linux commands](#) on page 48

Related information

[NetApp Knowledgebase Answer 1087383: Oracle RAC One Node database is skipped for performing SnapCenter operations](#)

Backing up Oracle database resource groups

A resource group is a collection of resources on a host or cluster. A backup operation on the resource group is performed on all resources defined in the resource group.

Before you begin



- You must have created a resource group with a policy attached.
- If you are backing up a resource that has a SnapMirror relationship to secondary storage, you must have SnapCenter admin privileges and the scadmin role must include the “snapmirror all” privilege.
- You must have assigned the aggregate that is being used by the backup operation to the storage virtual machine (SVM) used by the database.
- You should ensure that all data volumes and archive log volumes belonging to the database are protected if secondary protection is enabled for that database.

About this task

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 occur automatically according to the schedule.

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 **Resource Group** from the **View** list.

You can search the resource group either by entering the resource group name in the search box or by clicking , and then selecting the tag. You can then click  to close the filter pane.

3. In the **Resource Groups** page, select the resource group that you want to back up, and then click **Back up Now**.

Note: If you have a federated resource group with two databases and one of the database has datafile on non-NetApp storage, the backup operation is aborted even though the other database is on NetApp storage.

4. In the **Backup** page, perform the following steps:
 - a. If you have associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.
 - b. Click **Backup**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.

Related tasks

[Creating backup policies for Oracle databases](#) on page 17

[Creating resource groups and attaching policies for Oracle databases](#) on page 21

[Monitoring backup operations](#) on page 26

[Viewing Oracle database backups and clones in the Topology page](#) on page 28

[Backing up Oracle databases using Linux commands](#) on page 48

Related information







[NetApp Knowledgebase Answer 1087383: Oracle RAC One Node database is skipped for performing SnapCenter operations](#)

Monitoring backup operations


You can monitor the progress of different backup operations by using the SnapCenter Jobs 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 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. Optional: 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.

Note: 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.

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

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


Monitoring 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.

About this task

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.

Canceling the SnapCenter Plug-in for Oracle Database backup operations

You can cancel backup operations that are either running, queued, or non-responsive. When you cancel a backup operation, the SnapCenter Server stops the operation and removes all the Snapshot copies from the storage if the backup created is not registered with SnapCenter Server. If the backup is already registered with SnapCenter Server, it will not roll back the already created Snapshot copy even after the cancellation is triggered.

Before you begin


- You must be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel only the log or full backup operation that are queued or running.
- You cannot cancel the operation after the verification has started.
If you cancel the operation before verification, the operation is canceled, and the verification operation will not be performed.
- You cannot cancel the backup operation after the catalog operations has started.
- You can cancel a backup operation from either the Monitor page or the Activity pane.
- In addition to using the SnapCenter GUI, you can use CLI commands to cancel operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Step

1. Perform one of the following actions:

From the...	Action
Monitor page	<p>a. In the left navigation pane, click Monitor > Jobs.</p> <p>b. Select the operation and click Cancel Job.</p>
Activity pane	<p>a. After initiating the backup job, click  on the Activity pane to view the five most recent operations.</p> <p>b. Select the operation.</p> <p>c. In the Job Details page, click Cancel Job.</p>

Result

The operation is canceled, and the resource is reverted to the original state. If the operation you canceled is non-responsive in the canceling or running state, you should run the `Cancel-SmJob -JobID <int> -Force` to forcefully stop the backup operation.

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)




Viewing Oracle database backups and clones in the Topology page

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

About this task

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups 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.

- If you have upgraded from SnapCenter 1.1, the clones on the secondary (mirror or vault) are not displayed under Mirror copies or Vault copies in the Topology page.
All of the clones created using SnapCenter 1.1 are displayed under the Local copies in SnapCenter 3.0.

Note: 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.

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 and total number of log backups.

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


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, mount, unmount, rename, and delete operations.

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

- If you have selected a log backup, you can only perform rename, mount, unmount, and delete operations.
- If you have cataloged the backup using Oracle Recovery Manager (RMAN), you cannot rename those cataloged backups.

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

Mounting and unmounting database backups

You can mount a single or multiple data and log only backups if you want to access the files in the backup. You can either mount the backup to the same host where the backup was created or to a remote host having same type of Oracle and host configurations. If you have manually mounted the backups, you should manually unmount the backups after completing the operation. You can mount the database backup only once on a host for a specific database. While performing an operation, you can mount only a single backup.

Related tasks

[Mounting a database backup](#) on page 30

[Unmounting a database backup](#) on page 31

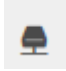
Mounting a database backup

You can manually mount a database backup if you want to access the files in the backup.

Before you begin

- If you have an Automatic Storage Management (ASM) database instance in NFS environment and want to mount the ASM backups, you must have added the ASM disk path `/var/opt/snapcenter/sco/backup_*/**/*/*/*` to the existing path defined in the `asm_diskstring` parameter.
- If you have an ASM database instance in NFS environment and want to mount the ASM log backups as part of recovery operation, you must have added the ASM disk path `/var/opt/snapcenter/scu/clones/**/*` to the existing path defined in the `asm_diskstring` parameter.
- If you want to mount to an alternate host you must ensure that the alternate host meets the following requirements:
 - Same UID and GID as that of the original host
 - Same Oracle version as that of original host
 - Same OS distribution and version as that of original host

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 **Database** or **Resource Group** from the **View** list.
3. Select the database either from the database details view or from the resource group details view.
The database topology page is displayed.
4. From the **Manage Copies** view, select **Backups** either from the primary or secondary (mirrored or replicated).
5. Select the backup from the table and click .

6. In the **Mount backups** page, select the host on which you want to mount the backup from the **Choose the host to mount the backup** drop-down list.

The mount path `/var/opt/snapcenter/sco/backup_mount/backup_name/database_name` is displayed.

If you are mounting the backup of an ASM database, the mount path `+diskgroupname_SID_backupid` is displayed.

7. Click **Mount**.

After mounting the backup, you can run the following command to retrieve the information related to the mounted backup:

```
./sccli Get-SmBackup -BackupName backup_name -ListMountInfo
```

If you have mounted an ASM database, you can run the following command to retrieve the information related to the mounted backup:

```
./sccli Get-Smbackup -BackupName diskgroupname_SID_backupid -listmountinfo
```

To retrieve the backup id run the following command:

```
./sccli Get-Smbackup -BackupName backup_name
```

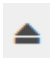
Unmounting a database backup

You can manually unmount a mounted database backup when you no longer want to access files on the backup.

Before you begin

You must have manually mounted a backup.

Steps

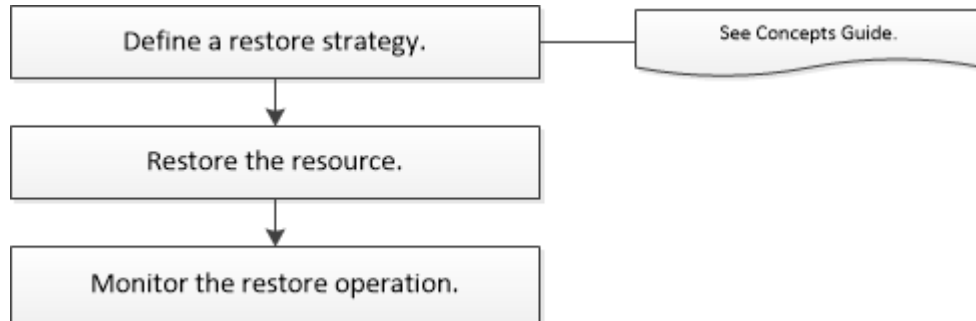
1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database either from the database details view or from the resource group details view.
The database topology page is displayed.
4. Select the backup that is mounted, and then click .
5. Click **OK**.

Restoring Oracle databases

The restore workflow includes planning, performing the restore operations, and monitoring the operations.

About this task

The following workflow shows the sequence in which you must perform the restore operation:



You can also use Linux commands manually or in scripts to perform restore operation. For detailed information about Linux commands, use the SnapCenter command help or see the command reference information.

Related tasks

[Restoring and recovering Oracle databases using Linux commands](#) on page 49

Related information

[Concepts](#)

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Restoring an Oracle database

In the event of data loss, you can use SnapCenter to restore data from one or more backups to your active file system and then recover the database. Recovery operation is performed by using archive logs of the database in an active file system.

Before you begin


- You should have defined your restore and recovery strategy.
- The SnapCenter administrator should have assigned you the storage virtual machines (SVMs) for both the source volumes and destination volumes if you are replicating Snapshot copies to a mirror or vault.
- If archive logs are pruned as part of backup, you should have manually mounted the required archive log backups.
- If you want to restore Oracle databases that are residing on a Virtual Machine Disk (VMDK), you should ensure that the guest machine has the required number of free slots for allocating cloned VMDKs.

- You should ensure that all data volumes and archive log volumes belonging to the database are protected if secondary protection is enabled for that database.
- You should ensure that the RAC One Node database is in nomount state to perform control file or full database restore.

About this task

When you restore a database, an operational lock file (`.sm_lock_dbsid`) is created on the Oracle database host in the `$ORACLE_HOME/dbs` directory to avoid multiple operations being executed on the database. After the database has been restored, the operational lock file is automatically removed.

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 **Database** or **Resource Group** from the **View** list.
3. Select the database from either the database details view or the resource group details view.
The database 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, perform the following tasks:
 - a. If you have selected a backup of a database in a Real Application Clusters (RAC) environment, select the RAC node.
 - b. Perform the following actions:

If you want to restore...	Do this...
All datafiles	Select All Datafiles .
Tablespaces	Select Tablespaces . You can specify the tablespaces that you want to restore.
Control files	Select Control files .
Redo log files	Select Redo log files . This option is available only for Data Guard standby or Active Data Guard standby databases.
Pluggable databases (PDBs)	Select Pluggable databases , and then specify the PDBs that you want to restore.
Pluggable database (PDB) tablespaces	Select Pluggable database (PDB) tablespaces , and then specify the PDB and the tablespaces of that PDB that you want to restore. This option is available only if you have selected a PDB for restore.

- c. Select **Change database state if needed for restore and recovery** to change the state of the database to the state required to perform restore and recovery operations.

The various states of a database from higher to lower are open, mounted, started, and shutdown. You must select this check box if the database is in a higher state but the state must be changed to a lower state to perform a restore operation. If the database is in a lower state

but the state must be changed to a higher state to perform the restore operation, the database state is changed automatically even if you do not select the check box.

Example

If a database is in the open state, and for restore the database needs to be in the mounted state, then the database state is changed only if you select this check box.

- d. Select **Force in place restore** if you want to perform in-place restore in the scenarios where new datafiles are added after backup or when LUNs are added, deleted, or re-created to an LVM disk group.

7. In the **Recovery Scope** page, perform the following actions:

You cannot perform restore with recovery from secondary backups if archive log volumes are not protected but data volumes are protected. You can restore only by selecting **No recovery**.

If you...	Do this...
Want to recover to the last transaction	Select All Logs .
Want to recover to a specific System Change Number (SCN)	Select Until SCN (System Change Number) .
Want to recover to a specific data and time	Select Date and Time . You must specify the date and time of the database host's time zone.
Do not want to recover	Select No recovery .
Want to specify any external archive log locations	Select Specify external archive log locations , and then specify the location of the external archive log files. If archive logs are pruned as part of backup, and you have manually mounted the required archive log backups, you must specify the mounted backup path as the external archive log location for recovery. NetApp Technical Report 4591: Database Data Protection: Backup, Recovery, Replication, and DR NetApp Knowledgebase Answer 1086191: SnapCenter Oracle restore or clone fails with the message 'ORA-00308: cannot open archived log / ORA_LOG/arch1_123_456789012.arc'

Note: Recovery is not supported for Data Guard standby and Active Data Guard standby databases.

8. In the **PreOps** page, enter the path and the arguments of the prescript that you want to run before the restore operation.

You must store the prescripts either in the `/var/opt/snapcenter/spl/scripts` path or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

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

9. In the **PostOps** page, perform the following steps:

- a. Enter the path and the arguments of the postscript that you want to run after the restore operation.

You must store the postscripts either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is

populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

- b. Select the check box if you want to open the database after recovery.

After restoring a container database (CDB) with or without control files, or after restoring only CDB control files, if you specify to open the database after recovery, then only the CDB is opened and not the pluggable databases (PDB) in that CDB.

In a RAC setup, only the RAC instance that is used for recovery is opened after recovery.

Note: After restoring a user tablespace with control files, a system tablespace with or without control files, or a PDB with or without control files, only the state of the PDB related to the restore operation is changed to the original state. The state of the other PDBs that were not used for restore are not changed to the original state because the state of those PDBs were not saved. You must manually change the state of the PDBs that were not used for restore.

10. 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**.

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

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

Related tasks

[Mounting a database backup](#) on page 30

Related information

[NetApp Knowledgebase Answer 1087383: Oracle RAC One Node database is skipped for performing SnapCenter operations](#)

[NetApp Knowledgebase Answer 1081874: SnapCenter with Oracle Plug-in Fails a Restore Operation with MarshallImpl.write Exception Error Due to Network Firewall Inactive Connection Setting](#)

[NetApp Knowledgebase Answer 1081879: SnapCenter with Oracle Plug-in Restore Operation Reverts to a Connect and Copy Operation, Instead of an In-Place Restore, When the Linux Hostname and Oracle RAC HOST_NAME Do Not Match](#)


Monitoring 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. Optional: In the **Jobs** page, perform the following steps:
 - a. Click  to filter the list so that only restore operations are listed.
 - b. Optional: 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 are completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. Optional: In the **Job Details** page, click **View logs**.
The **View logs** button displays the detailed logs for the selected operation.

Canceling restore operations

You can cancel restore jobs that are queued.

Before you begin


- You must be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel a restore operation from either the Monitor page or the Activity pane.
- You cannot cancel a running restore operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the restore operations.
- The **Cancel Job** button is disabled for restore operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Step

1. Perform one of the following actions:

From the...	Action
Monitor page	<p>a. In the left navigation pane, click Monitor > Jobs.</p> <p>b. Select the job and click Cancel Job.</p>
Activity pane	<p>a. After initiating the restore operation, click  on the Activity pane to view the five most recent operations.</p> <p>b. Select the operation.</p> <p>c. In the Job Details page, click Cancel Job.</p>

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Cloning Oracle database

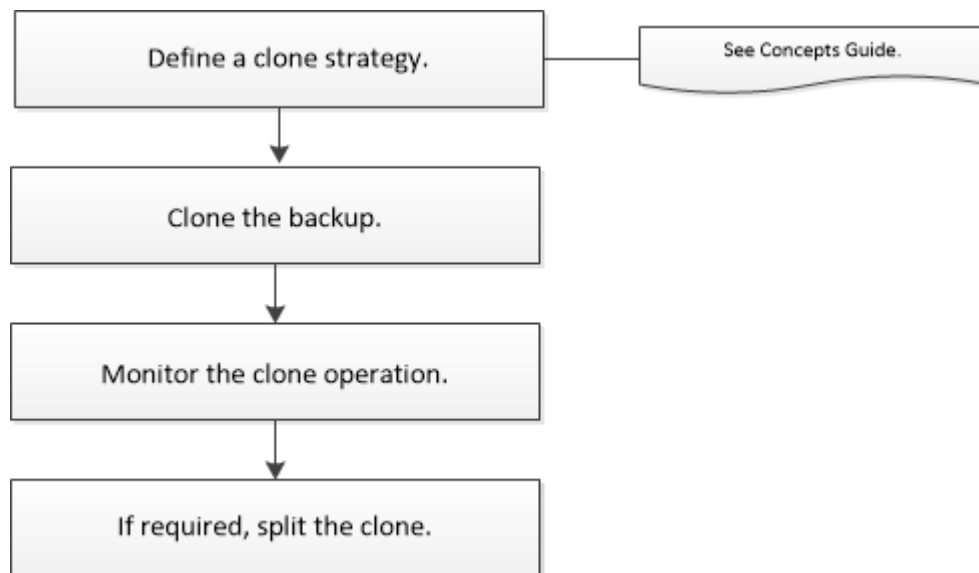
The clone workflow includes planning, performing the clone operation, and monitoring the operation.

About this task

You might clone databases for the following reasons:

- To test functionality that has to be implemented using the current database structure and content during application development cycles.
- To populate data warehouses using data extraction and manipulation tools.
- To recover data that was mistakenly deleted or changed.

The following workflow shows the sequence in which you must perform the clone operation:



You can also use Linux commands to perform clone operation. For detailed information about Linux commands, use the SnapCenter command help or see the command reference information.

Related tasks

[Cloning Oracle database backups using Linux commands](#) on page 50

[Refreshing a clone](#) on page 50

Related information

[Concepts](#)

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Cloning an Oracle database

You can use SnapCenter to clone a database using the backup of the database. The clone operation creates a copy of the database data files, and creates new online redo log files and control files. The database can be optionally recovered to a specified time, based on the specified recovery options.

Before you begin

- You should have created a backup of the database using SnapCenter.
The backups should be successfully created. You should have either online data and log backups or offline (mount or shutdown) backups created for the cloning operation to succeed.
- If you want to customize the control file or redo log file paths, you should have preprovisioned the required file system or Automatic Storage Management (ASM) disk group.
By default, redo log and control files of the cloned database are created on the ASM disk group or the file system provisioned by SnapCenter for the data files of the clone database.
- The clone can be created on the same host as that of the source database or on an alternate host. If you are creating the clone on an alternate host, the alternate host must meet the following requirements:
 - SnapCenter Plug-in for Oracle Database should be installed on the alternate host.
 - The clone host should be able to discover LUNs from primary or secondary storage.
 - If you are cloning from primary storage or secondary (Vault or Mirror) storage to an alternate host, then make sure that an iSCSI session is either established between the secondary storage and the alternate host, or zoned properly for FC.
[Linux Host Utilities Installation and Setup Guide](#).
 - If you are cloning from Vault or Mirror storage to the same host, then make sure that an iSCSI session is either established between the Vault or Mirror storage and the host, or zoned properly for FC.
 - If you are cloning in a virtualized environment, ensure that an iSCSI session is either established between the primary or secondary storage and ESX server hosting the alternate host, or zoned properly for FC.
 - If the source database is an ASM database:
 - The ASM instance should be up and running on the host where the clone will be performed.
 - The ASM disk group should be provisioned prior to the clone operation if you want to place archive log files of the cloned database in a dedicated ASM disk group.
 - The name of the data disk group can be configured, but ensure that the name is not used by any other ASM disk group on the host where the clone will be performed.
Data files residing on the ASM disk group are provisioned as part of the SnapCenter clone workflow.
- The protection type for the data LUN and the log LUN, such as mirror, vault, or mirror-vault, should be the same to discover secondary locators during cloning to an alternate host using log backups.
- You should set the value of `exclude_seed_cdb_view` to **FALSE** in the source database parameter file to retrieve seed PDB related information for cloning a backup of 12c database. The seed PDB is a system-supplied template that the CDB can use to create PDBs. The seed PDB is named `PDB$SEED`. For information about PDB\$SEED, see the Oracle Doc ID 1940806.1.

Note: You should set the value before backing up 12c database.

- SnapCenter supports backup of file systems that are managed by the autofs subsystem. If you are cloning the database, ensure that data mount points are not under the root of the autofs mount point because the root user of the plug-in host does not have permission to create directories under the root of the autofs mount point.

If control and redo log files are under data mount point, you should modify the control file path and redo log file path accordingly.

Note: You can manually register the new cloned mount points with the autofs subsystem. The new cloned mount points will not be registered automatically.

- If you have a TDE (auto login) and want to clone the database on the same or alternate host, you should copy wallet (key files) under `/etc/ORACLE/WALLET/$ORACLE_SID` from the source database to the cloned database.

About this task

SnapCenter creates a stand-alone database when cloned from an Oracle RAC database backup. SnapCenter supports creating clone from the backup of a Data Guard standby and Active Data Guard standby databases.

During cloning, SnapCenter mounts the log backup for recovery operations. After recovery, the log backup is unmounted. All such clones are mounted under `/var/opt/snapcenter/scu/clones/`. If you are using ASM over NFS, you should add `/var/opt/snapcenter/scu/clones/*/*` to the existing path defined in the `asm_diskstring` parameter.


While cloning a backup of an ASM database in a SAN environment, udev rules for the cloned host devices are created at `/etc/udev/rules.d/999-scu-netapp.rules`. These udev rules associated with the cloned host devices are deleted when you delete the clone.

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 **Database** or **Resource Group** from the **View** list.
3. Select the database either from the database details view or from the resource group details view.

The database 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 Data backup from the table, and then click .

6. In the **Name** page, enter the SID of the clone.

The clone SID is not available by default, and the maximum length of the SID is 8 characters.

Note: You should ensure that no database with the same SID exists on the host where the clone will be created.

7. In the **Locations** page, perform the following actions:

For this field...	Do this...
Clone host	By default, the source database host is populated. If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host.

For this field...	Do this...
Datafile locations	<p>By default, the datafile location is populated.</p> <p>The SnapCenter default naming convention for SAN or NFS file systems is <i>FileSystemNameofsourcedatabase_CLONESID</i>.</p> <p>The SnapCenter default naming convention for ASM disk groups is <i>SC_HASHCODEofDISKGROUP_CLONESID</i>. The <i>HASHCODEofDISKGROUP</i> is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group.</p> <p>Note: If you are customizing the ASM disk group name, ensure that the name length adheres to the maximum length supported by Oracle.</p> <p>If you want to specify a different path, you must enter the datafile mount points or ASM disk group names for clone database. When you customize the datafile path, you must also change the control file and redo log file ASM disk group names or file system either to the same name used for data files or to an existing ASM disk groups or file system.</p>
Control files	<p>By default, the control file path is populated.</p> <p>The control files are placed in the same ASM disk group or file system as that of the data files. If you want to override the control file path, you can provide a different control file path.</p> <p>Note: The file system or the ASM disk group should exist on the host.</p> <p>By default, the number of control files will be same as that of the source database. You can modify the number of control files but a minimum of one control file is required to clone the database.</p> <p>You can customize the control file path to a different file system (existing) than that of the source database.</p>

For this field...	Do this...
Redo logs	<p>By default, the redo log file group, path, and their sizes are populated.</p> <p>The redo logs are placed in the same ASM disk group or file system as that of the data files of the cloned database. If you want to override the redo log file path, you can customize the redo log file path to a different file system than that of the source database..</p> <p>Note: The new file system or the ASM disk group should exist on the host.</p> <p>By default, the number of redo log groups, redo log files, and their sizes will be same as that of the source database. You can modify the following parameters:</p> <ul style="list-style-type: none"> number of redo log groups <p>Note: A minimum of three redo log groups are required to clone the database.</p> redo log files in each group and their path <p>Note: A minimum of one redo log file is required in the redo log group to clone the database.</p> <p>You can customize the redo log file path to a different file system (existing) than that of the source database.</p> sizes of the redo log file

8. In the **Credentials** page, perform the following actions:

For this field...	Do this...
Run As name for sys user	<p>Select the Run As account to be used for defining the sys user password of the clone database.</p> <p>If <code>SQLNET.AUTHENTICATION_SERVICES</code> is set to NONE in <code>sqlnet.ora</code> file on the target host, you should not select None as the Run As account in the SnapCenter GUI.</p>
ASM Instance Run As name	<p>Select None if OS authentication is enabled for connecting to the ASM instance on the clone host.</p> <p>Otherwise, select the Oracle ASM Run As configured with either “sys” user or an user having “sysasm” privilege applicable to the clone host.</p>

The Oracle home, user name, and group details are automatically populated from the source database. You can change the values based on the Oracle environment of the host where the clone will be created.

Database Credentials for the clone

Run As name for sys user + ⓘ

ASM instance Run As name + ⓘ

ASM Port

Oracle Home Settings ⓘ

Oracle Home

Oracle OS User


Oracle OS Group

9. In the **PreOps** page, perform the following steps:

- a. Enter the path and the arguments of the prescript that you want to run before the clone operation.

You must store the prescript either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.

- b. In the **Database Parameter settings** section, modify the values of prepopulated database parameters that are used to initialize the database.

You can add additional parameters by clicking .

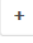
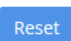
Note: Fast recovery area (FRA) is not defined is the prepopulated database parameters. You can configure FRA by adding the related parameters.

Note: The default value of `log_archive_dest_1` is `$ORACLE_HOME/clone_sid` and the archive logs of the cloned database will be created in this location. If you have deleted the `log_archive_dest_1` parameter, the archive log location is determined by Oracle. You can define a new location for archive log by editing `log_archive_dest_1` but ensure that the file system or disk group should be existing and made available on the host.

[NetApp Knowledgebase Answer 1086191: SnapCenter Oracle restore or clone fails with the message 'ORA-00308: cannot open archived log /ORA_LOG/arch1_123_456789012.arc'](#)

Database Parameter settings

audit_file_dest	/ora01/app/oracle/admin/tst/adump	×	⬆
audit_trail	DB	×	⬆
log_archive_dest_1	LOCATION=/ora01/app/oracle/product/12c/db...	×	⬆
log_archive_format	%t_%s_%r.dbf	×	⬇

Click **Reset** to get the default database parameter settings.

- 10.** In the **PostOps** page, **Recover database** and **Until Cancel** are selected by default to perform recovery of the cloned database.

SnapCenter performs recovery by mounting the latest log backup that have the unbroken sequence of archive logs after the data backup that was selected for cloning. The log and data backup should be on primary storage to perform the clone on primary storage and log and data backup should be on secondary storage to perform the clone on secondary storage.

The **Recover database** and **Until Cancel** options are not selected if SnapCenter fails to find the appropriate log backups. You can provide the external archive log location if log backup is not available in **Specify external archive log locations**. You can specify multiple log locations.

Note: If you want to clone a source database that is configured to support flash recovery area (FRA) and Oracle Managed Files (OMF), the log destination for recovery must also adhere to OMF directory structure.

The PostOps page is not displayed if the source database is a Data Guard standby or an Active Data Guard standby database. For Data Guard standby or an Active Data Guard standby database, SnapCenter does not provide an option to select the type of recovery in the SnapCenter GUI but the database is recovered using Until Cancel recovery type without applying any logs.

Filed name	Description
Until Cancel	SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning. The cloned database is recovered till the missing or corrupt log file.
Date and time	SnapCenter recovers the database up to a specified date and time. The accepted format is <i>mm/dd/yyyy hh:mm:ss</i> . Note: The time can be specified in 24 hour format.
Until SCN (System Change Number)	SnapCenter recovers the database up to a specified system change number (SCN).
Specify external archive log locations	Specify the external archive log location.
Create new DBID	By default Create new DBID check box is selected to generate a unique number (DBID) for the cloned database differentiating it from the source database. Clear the check box if you are fine to have the same DBID for the cloned database as that of the source database. In this scenario, if you want to register the cloned database with the external RMAN catalog where the source database is already registered, the operation fails.

11. 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 restore operation performed, select **Attach Job Report**.

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

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

Note: While performing recovery as part of clone create operation, even if recovery fails, the clone is created with a warning. You can perform manual recovery on this clone to bring the clone database to consistent state.

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

Result

After cloning the database you can refresh the resources page to list the cloned database as one of the resource available for backup. The cloned database can be protected like any other database using the standard backup workflow or can be included in a resource group (either newly created or existing). The cloned database can be further cloned (clone of clones).

Note: If you have not performed recovery while cloning, the backing up of the cloned database might fail due to improper recovery and you might have to perform manual recovery. The log backup can also fail if default location which was populated for archive logs is on a non-NetApp storage or if the storage system is not configured with SnapCenter.

Related tasks







[Viewing Oracle database backups and clones in the Topology page](#) on page 28

Monitoring clone operations in SnapCenter


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**.

Canceling clone operations

You can cancel clone operations that are queued.

Before you begin


- You must be logged in as the SnapCenter Admin or job owner to cancel operations.

About this task

- You can cancel a clone operation from either the Monitor page or the Activity pane.
- You cannot cancel a running clone operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the clone operations.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Step

1. Perform one of the following actions:

From the...	Action
Monitor page	<ol style="list-style-type: none"> a. In the left navigation pane, click Monitor > Jobs. b. Select the operation, and click Cancel Job.
Activity pane	<ol style="list-style-type: none"> a. After initiating the clone operation, click  on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click Cancel Job.

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Splitting an Oracle Database 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 of the clone are deleted.
- For information about clone split operation limitations, see the *Logical Storage Management Guide*.
[ONTAP 9 Logical Storage Management Guide](#)

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 **Database** from the **View** list.
3. Select the cloned resource, (for example, the database or LUN) and then click .
4. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.

After you finish

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" />
```

Backing up, restoring, and cloning using Linux commands

The SnapCenter Plug-in for Oracle Database includes Linux commands for scripting of backup, restore, recovery, and clone operations.

The following are common tasks you might perform using Linux commands:

- Backing up Oracle databases
- Restoring and recovering Oracle databases
- Cloning Oracle database backups

For detailed information about Linux commands, use the SnapCenter command help or see the command reference information.

[*SnapCenter Software 4.1 Linux Command Reference Guide*](#)

Backing up Oracle databases using Linux commands

The backup workflow includes planning, identifying the resources for backup, creating backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

Before you begin

- You must have added the storage system connections and created the Run As account using the commands `Add-SmStorageConnection` and `Add-SmRunAs`.
- You must have established the connection session with the SnapCenter Server using the command `Open-SmConnection`.

You can have only one SnapCenter account login session and the token is stored in the Linux user home directory.

Note: The connection session is valid only for 24 hours. However, you can create a token with the `TokenNeverExpires` option to create a token that never expires and session will always be valid.

About this task

You must execute the following commands to establish the connection with the SnapCenter Server, discover the Oracle database instances, add policy and resource group, backup and verify the backup.

For detailed information on Linux commands, use the SnapCenter command help or see the command reference information.

[*SnapCenter Software 4.1 Linux Command Reference Guide*](#)

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: `Open-SmConnection`
2. Perform host resources discovery operation: `Get-SmResources`
3. Configure Oracle database credentials and preferred nodes for backup operation of a Real Application Cluster (RAC) database: `Configure-SmOracleDatabase`

4. Create a backup policy: `Add-SmPolicy`
5. Retrieve the information about the secondary (SnapVault or SnapMirror) storage location : `Get-SmSecondaryDetails`

This command retrieves the primary to secondary storage mapping details of a specified resource. You can use the mapping details to configure the secondary verification settings while creating a backup resource group.
6. Add a resource group to SnapCenter: `Add-SmResourceGroup`
7. Create a backup: `New-SmBackup`.

You can poll the job using the `WaitForCompletion` option. If this option is specified, then the command continues to poll the server until the completion of the backup job.
8. Retrieve the logs from SnapCenter: `Get-SmLogs`

Restoring and recovering Oracle databases using Linux commands

The restore and recovery workflow includes planning, performing the restore and recovery operations, and monitoring the operations.

Before you begin

- You must have established the connection session with the SnapCenter Server.

About this task

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

For detailed information on Linux commands, use the SnapCenter command help or see the command reference information.

[*SnapCenter Software 4.1 Linux Command Reference Guide*](#)

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: `Open-SmConnection`
2. Retrieve the information about the backups that you want to restore: `Get-SmBackup`
3. Retrieve the detailed information about the specified backup: `Get-SmBackupDetails`

This command retrieves the detailed information about the backup of a specified resource with a given backup ID. The information includes database name, version, home, start and end SCN, tablespaces, pluggable databases, and its tablespaces.
4. Restore data from the backup: `Restore-SmBackup`

Cloning Oracle database backups using Linux commands

The clone workflow includes planning, performing the clone operation, and monitoring the operation.

Before you begin

- You must have established the connection session with the SnapCenter Server.

About this task

You must execute the following commands to create the Oracle database clone specification file and initiate the clone operation.

For detailed information on Linux commands, use the SnapCenter command help or see the command reference information.

[*SnapCenter Software 4.1 Linux Command Reference Guide*](#)

Steps

1. Create an Oracle database clone specification from a specified backup: `New-SmOracleCloneSpecification`

Note: If secondary data protection policy is unified mirror-vault, then specify only `-IncludeSecondaryDetails`. You do not have to specify `-SecondaryStorageType`.

This command automatically creates an Oracle database clone specification file for the specified source database and its backup. You must also provide a clone database SID so that the specification file created has the automatically generated values for the clone database which you will be creating.

Note: The clone specification file is created at `/var/opt/snapcenter/sco/clone_specs`.

2. Initiate a clone operation from a clone resource group or an existing backup: `New-SmClone`

This command initiates a clone operation. You must also provide an Oracle clone specification file path for the clone operation. You can also specify the recovery options, host where the clone operation to be performed, prescripts, postscripts, and other details.

By default, the archive log destination file for the clone database is automatically populated at `$ORACLE_HOME/CLONE_SIDs`.

Refreshing a clone

You can refresh the clone by running the “Refresh-SmClone” command. This command creates a backup of the database, deletes the existing clone, and creates a clone with the same name.

About this task

You should be aware of some of best practices that you should follow before performing this operation.

- Create an online full backup or an offline data backup policy with no scheduled backups enabled.
- Configure the email notification in the policy for backup failures only.
- Define the retention count for the on-demand backups appropriately to ensure that there are no unwanted backups.

- Ensure that only an online full backup or an offline data backup policy is associated with resource group which is identified for refresh clone operation.
- Create a resource group with only one database.
- If a cron job is created for the clone refresh command, ensure that the SnapCenter schedules and the cron schedules are not overlapping for the database resource group.
For a cron job created for the clone refresh command, ensure that you run `Open-SmConnection` after every 24hrs.
- Ensure that the clone SID is unique for a host.
If multiple refresh clone operations use the same clone specification file or use the clone specification file with same clone SID, existing clone with the SID on the host will be deleted and then the clone will be created.
- Ensure that the backup policy is enabled with secondary protection and the clone specification file is created with “`-IncludeSecondaryDetails`” to create the clones using secondary backups.
 - If the primary clone specification file is specified but the policy has secondary update option selected, the backup will be created, and update will get transferred to secondary. However, the clone will be created from the primary backup.
 - If the primary clone specification file is specified and the policy does not have secondary update option selected, the backup will be created on primary and clone will be created from primary.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: `Open-SmConnection`
2. Create an Oracle database clone specification from a specified backup: `New-SmOracleCloneSpecification`

Note: If secondary data protection policy is unified mirror-vault, then specify only `-IncludeSecondaryDetails`. You do not have to specify `-SecondaryStorageType`.

This command automatically creates an Oracle database clone specification file for the specified source database and its backup. You must also provide a clone database SID so that the specification file created has the automatically generated values for the clone database which you will be creating.

Note: The clone specification file is created at `/var/opt/snapcenter/sco/clone_specs`.

3. Run `Refresh-SmClone`.

If the operation fails with the PL-SCO-20032: `canExecute operation failed with error: PL-SCO-30031: Redo log file +SC_2959770772_clmdb/clmdb/redolog/redo01_01.log exists` error messages, specify a higher value for “`-WaitToTriggerClone`”

For detailed information on Linux commands, use the SnapCenter command help or see the command reference information.

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Troubleshooting data protection operations

If you encounter unexpected behavior while performing data protection operations, you can use the log files to identify the cause and to resolve the problem.

The log files are located at `/var/opt/snapcenter/spl/logs`. You can also download the log files from the SnapCenter user interface by clicking **Monitor > Logs > Download**.

Discovery operation takes long time to complete

Description

Discovery operation takes long time to complete if the storage is full, if the database is not reachable, or if any of the volumes are offline on which the database is residing. Discovery operation will be aborted only when it exceeds the discovery timeout value.

Corrective action

You must fix the issue that is causing the discovery operation to run for long time.

Unable to add Linux host to SnapCenter

Description

When you try to add a Linux host, the operation might fail. This issue occurs if the Oracle Java Database Connectivity (JDBC) port or SnapCenter Plug-in Loader (SPL) port is not free. The default JDBC port is 27216 and the default SPL port is 8145.

Error message

```
Unexpected failure of plug-ins discovery operation:
java.lang.IllegalStateException: Address already in use (Bind failed).
```

Corrective action

- If the default JDBC port is not free, perform the following steps:
 1. Assign a new port number to the `remote.registry.ocijdbc.port` parameter in the `/var/opt/snapcenter/sco/etc/sco.properties` file.
 2. Restart the SPL service by running the following command:


```
/opt/NetApp/snapcenter/spl/bin/spl restart
```
- If the SPL port is not free, you must specify a free port for SPL while adding the host from the SnapCenter UI.

Scanning of host bus adapters takes long time to complete

Description

Scanning of host bus adapters (HBAs) takes long time to complete because SnapCenter scans all the host bus adaptors present in the host.

Corrective action

You must perform the following steps:

1. Edit the `LinuxConfig.pm` file located at `/opt/NetApp/snapcenter/spl/plugins/scu/scucore/modules/SCU/Config`.
2. Add the missing HBA names of the host to the `HBA_DRIVER_NAMES` parameter.
3. Set the value of the `SCSI_HOSTS_OPTIMIZED_RESCAN` parameter to `1` to rescan only those HBAs that are listed in `HBA_DRIVER_NAMES`.

Backup fails during the discovery of file system on a VM

Description

When you perform database backup which is provisioned on a VMDK file system, backup operation fails while discovering the file system.

Error message

Failed to retrieve the unit serial number for the device.

Corrective action

Log into vCenter and navigate to **VM options > Advanced > Edit configuration** to set the value of `disk.enableUUID` to `true` for the VM.

Backup operation fails during the storage discovery process

Description

The backup operation fails during the storage discovery process when the storage system is running on Data ONTAP operating in 7-Mode.

Error message

Failed - There is no storage connection configured for the user administrator.

Corrective action

You must be sure to use the supported storage environment. SnapCenter supports only clustered ONTAP storage systems.

Backup operation fails if database query is timed out

Description

Backup operation fails if database query execution time exceeds the timeout value either because of offline volume, database is not reachable, archive log destination is full, or storage volume is full.

Error message

Failed to update metadata due to: ORACLE-10012: Error executing SQL "
SELECT TS.NAME as TS_NAME, DBATS.CONTENTTS as DBATS_CONTENTTS, DBATS.STATUS

```
as DBATS_STATUS, DBATS.BLOCK_SIZE as DBATS_BLOCK_SIZE, TF.FILE_NAME AS
TF_NAME, TF.BYTES as TF_BYTES, V$TABLESPACE TS WHERE TF.TABLESPACE_NAME =
TS.NAME AND TS.NAME = DBATS.TABLESPACE_NAME ORDER BY TS.NAME" within 2100
seconds against Oracle database v2003d7. Check the database status or
increase timeout.
```

ORA-01013: user requested cancel of current operation.

If the SQL query exceeds the timeout value, the SQL query operation is cancelled.

Corrective action

You must fix the scenario that caused the operation to fail.

If the database queries are slow, change the value of the ORACLE_SQL_QUERY_TIMEOUT and ORACLE_PLUGIN_SQL_QUERY_TIMEOUT parameters by running the following command:

```
/opt/Netapp/snapcenter/spl/bin/sccli Set-SmConfigSettings -
ConfigSettingsType Plugin -PluginCode SCO -ConfigSettings
"KEY=ORACLE_SQL_QUERY_TIMEOUT,VALUE=user_defined_value" -ConfigSettings
"KEY=ORACLE_PLUGIN_SQL_QUERY_TIMEOUT=user_defined_value"
```

After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command:

```
/opt/NetApp/snapcenter/spl/bin/spl restart
```

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Backup operation might fail when the OS group ID of the Oracle database administrator is changed

Description

Backup operations might fail if you have changed the OS group ID of the Oracle database administrator.

Error message

Failed to collect metadata for control file.

Corrective action

You must change the OS group ID of the Oracle database administrator to the OS group ID that was used when the last successful backup was created.

Backup fails with error: Storage system(s) may need to be added, also ensure that the associated host is in a connected state

Description

A backup failed because the preferred IP address that was configured for the SVM went down. When the preferred IP comes up again, the SnapCenter cache is not automatically refreshed. Therefore, SnapCenter could not find the Preferred IP when attempting to perform the backup.

Corrective action

Refresh the SnapCenter cache for the SVM:

1. In the left navigation pane of the SnapCenter GUI, click **Storage Systems**.
2. In the Storage Systems page, select the storage system used by the backup, and then click **Modify**.
3. Make sure that the **Preferred IP** check box is selected and that the IP is correct.
4. Reenter the storage system password, and then click **OK**.
This action refreshes the SnapCenter cache and updates the storage system configuration.

Backup operation might fail if external RMAN catalog database has issues

Description

The backup operation fails in the following scenarios:

- The RMAN catalog database is down.
- The credentials for logging in to the RMAN catalog database is incorrect.
- The RMAN catalog database version and target database version is incompatible.
- The metadata in a recovery catalog schema does not match with the metadata in a target database control file.

Corrective action

Resolve the issues specific to the RMAN catalog database.

Detach the RMAN Run As credentials configured for the target database in SnapCenter until the RMAN catalog database issues are resolved.

Cataloging and uncataloging with Oracle RMAN will fail if the execution time is beyond the timeout value

Description

Catalog and uncatalog operations fail if the operation time exceeds the time out value specified for the ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT parameter.

Corrective action

You must modify the value of the ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT parameter by running the following command:

```
/opt/Netapp/snapcenter/spl/bin/sccli Set-SmConfigSettings -
ConfigSettingsType Plugin -PluginCode SCO -ConfigSettings
"KEY=ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT,VALUE=user_defined_value"
```

After modifying the value of the parameter, restart the SnapCenter Plug-in Loader (SPL) service by running the following command:

```
/opt/NetApp/snapcenter/spl/bin/spl restart
```

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Unable to find Snapshot copy after successfully creating the backup

Description

Backup operations with an update to SnapVault or SnapMirror might fail.

Depending on the rate of change of data between Snapshot copies, the time taken to update the Snapshot copy to secondary varies. If the delta between Snapshot copies is very high, it is recommended to change SnapshotCheckRetry and SnapshotCheckTimeout values.

Error message

Snapshot copy could not be found on the destination storage system.

Corrective action

You must include the following parameters and specify the value in the appsetting section of the SMCoreServiceHost.exe.Config file located under SmCore in the SnapCenter Server.

- `<add key="SnapshotCheckRetry" value=retry_value/>`
The `retry_value` assigned to `SnapshotCheckRetry` defines the maximum number of retries that are performed to discover the Snapshot copies on the secondary location.
- `<add key="SnapshotCheckTimeout" value=timeout_value/>`
The `timeout_value` (milliseconds) assigned to `SnapshotCheckTimeout` defines the wait or sleep period for every retry.

Backup operation fails if Snapshot copies on the secondary storage reaches maximum limit

Description

When the number of Snapshot copies on the secondary storage (mirror-vault) reaches the maximum limit, the activity to register backup and apply retention in the backup operation fails.

Error message

This Snapshot copy is currently used as a reference Snapshot copy by one or more SnapMirror relationships. Deleting the Snapshot copy can cause future SnapMirror operations to fail.

Corrective action

You should configure SnapMirror retention policy for the secondary storage to avoid reaching the maximum limit of Snapshot copies on the secondary storage.

Unable to update the SnapMirror and SnapVault status

Description

Mirror and Vault backup copies are not listed on the topology page even if data and log volumes are successfully protected. The issue occurs if the value assigned to SnapmirrorStatusUpdateWaitTime is less.

Corrective action

You should increase the value assigned to SnapmirrorStatusUpdateWaitTime using Set-SmConfigSettings Powershell cmdlet.

For example,

```
Set-SmConfigSettings -Server-configSettings
@{"SnapmirrorStatusUpdateWaitTime"=120000}
```

ASM backup verification fails

Description

Backup verification will fail if OS authentication is disabled on the Automatic Storage Management (ASM) host and the ASM database is configured for DB authentication in SnapCenter.

Corrective action

You must enable OS authentication on the ASM host.

Backup verification fails when files are not accessible

Description

Backup verification fails with the error code DBV-00100 specified file if the file is not accessible and the mount point is unavailable during the verification process.

Error message

Backup verification failed.

Corrective action

Modify values of the VERIFICATION_DELAY and VERIFICATION_RETRY_COUNT parameters in sco.properties.

Note: The VERIFICATION_DELAY parameter specifies the number of seconds to wait for completing the verification process, and VERIFICATION_RETRY_COUNT parameter specifies the number of time verification operation can be retried.

After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command:

```
/opt/NetApp/snapcenter/spl/bin/spl restart
```

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Backup verification is timed out

Description

For large databases (size in TBs), the verification operation might time out.

Corrective action

You should increase the timeout values.

- Increase the value of RESTTimeout to **86400000** seconds in C:\Program Files\NetApp\SMCore\SMCoreServiceHost.exe.config file in the SnapCenter Server host.
While modifying the values, ensure that there are no running jobs.
After increasing the value, you should restart the SnapCenter SMCore service.
- Increase the SnapCenter Server RESTTimeout to **86400000** seconds by running the following commands:
 - `sccli Get-SmConfigSettings -ConfigSettingsType Server -Key RESTTimeout`
The current timeout value is displayed.
 - `sccli Set-SmConfigSettings -ConfigSettingsType Server -ConfigSettings "KEY=RESTTimeout,VALUE=86400000"`

Disk paths are not included in the `asm_diskstring` database parameter

Description

By default, the `ASM_DISKSTRING_UPDATE` parameter is set to **false** in the `sco.properties` file. This parameter is set to false assuming that the value assigned to `asm_diskstring` includes the cloned disks path as well. However, sometimes the value assigned might not include the cloned disks path.

Corrective action

You must set the value of the `ASM_DISKSTRING_UPDATE` parameter to **true** to update the `asm_diskstring` database parameter to include the cloned disks path. After setting the `ASM_DISKSTRING_UPDATE` parameter to **true**, you must restart the SnapCenter Plug-in Loader (SPL) service.

Related information

[Installing and setting up SnapCenter](#)

Failed to mount ASM log backups as part of recovery operation

Description

If you have an ASM database instance in NFS environment, mounting of ASM log backups as part of recovery operation might fail if the appropriate ASM disk path is not defined in the `asm_diskstring` parameter.

Error message

```
ASM-00015: Mounting of ASM Disk Group <ASM_DISKGROUP_NAME> failed:
ORACLE-10003: Error executing SQL "ALTER DISKGROUP <ASM_DISKGROUP_NAME>
MOUNT RESTRICTED" against Oracle database +ASM: ORA-15032: not all
alterations performed ORA-15017: diskgroup "<ASM_DISKGROUP_NAME>" cannot be
mounted ORA-15040: diskgroup is incomplete
```

Corrective action

You should add the ASM disk path `/var/opt/snapcenter/scu/clones/*/*` to the existing path defined in the `asm_diskstring` parameter.

Unable to change the database state from shutdown to mount

Description

After creating an offline backup of a standalone Oracle 12c Automatic Storage Management (ASM) database, SnapCenter fails to change the state of Oracle database from shutdown to mount.

Error message

Resource failed with error PL-SCO-20005: Unquiescing of database failed with error ORACLE-20001: Error trying to change state to MOUNTED for database instance *database_instance*.

Corrective action

The database state change fails because of an Oracle issue in the Oracle 12.1.0.2 standalone ASM configuration (Oracle bug 18894342). You must apply the Oracle patch 18894342. For information about this Oracle issue, see the Oracle Doc ID 1922908.1.

Restore operation of datafiles and control files fail

Description

If you have disabled OS authentication and enabled Oracle database authentication for an Oracle database, and when you try to perform a restore of datafiles and control files of that database, the operation fails.

Corrective action

You must configure the static listener in the `listener.ora` file available at `$ORACLE_HOME/network/admin` and then retry the operation.

Restore from a secondary SnapMirror or SnapVault location fails

Description

Restore operation from a secondary SnapMirror or SnapVault location fails when load-sharing mirror (LSM) is configured on the primary volume. This issue occurs if you are using Data ONTAP 8.3 or later.

Error message

Destination *dest_vol* cannot be the source or destination of a load-sharing relationship.

Corrective action

You can perform one of the following:

- Specify a high retention count to retain more number of backup copies on the primary storage so that the restore operation can be performed from the primary storage.
- Mount the backup from the secondary storage and manually copy the files that have to be restored.

Restore operation might fail if the database size is in terabytes

Description

If the database size is in terabytes (TB), the restore operation might fail when the default timeout values are used.

Corrective action

You must increase the value of the `SCORestoreTimeout` parameter by running the `Set-SmConfigSettings` command. The default timeout value is 3 days (4320 minutes). For example, if you want to change timeout value to 4 days, you must run:

```
sccli Set-SmConfigSettings -ConfigSettingsType Server -ConfigSettings
"KEY=SCORestoreTimeout,VALUE=5760"
```

You can also run the `Get-SmConfigSettings` command to view the value of the `SCORestoreTimeout` parameter. For example:

```
sccli Get-SmConfigSettings -ConfigSettingsType Server -Key
SCORestoreTimeout -ShowDescription
```

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Restore operation fails when you select a backup of an orphan incarnation

Description

The restore operation fails either because the specified datafile was restored from a backup that was taken during a period of time that has already been discarded by a `RESETLOGS` operation or because Oracle cannot identify the database incarnation of the datafile.

Error message

```
ORA-19909: datafile %s belongs to an orphan incarnation.
```

Corrective action

Restore the datafile from a backup that belongs to either the current incarnation or a prior incarnation of the database.

The System Change Number (SCN) of the backup is displayed in the last column of the Topology page. You can use any of the backups with an SCN that is less than the current incarnation number. You should run the `list incarnation` command from the `RMAN` prompt of the target database to identify the current incarnation number.

In the following example, an earlier incarnation is changed to the current incarnation and all of the incarnations after that have become orphan incarnations. You can select any backup with an SCN that is less than or equal to 11061559.

```

RMAN> list incarnation;
List of Database Incarnations
DB Key   Inc Key DB Name   DB ID           STATUS   Reset SCN   Reset Time
-----
1        1       PDB_251   913203095      PARENT   1594143     26-SEP-17
2        2       PDB_251   913203095      PARENT   11031188    15-JAN-18
3        3       PDB_251   913203095      ORPHAN   11032983    15-JAN-18
4        4       PDB_251   913203095      PARENT   11032983    16-JAN-18
5        5       PDB_251   913203095      PARENT   11056390    16-JAN-18
8        8       PDB_251   913203095      ORPHAN   11059991    16-JAN-18
9        9       PDB_251   913203095      PARENT   11059991    16-JAN-18
6        6       PDB_251   913203095      ORPHAN   11060900    16-JAN-18
13       13      PDB_251   913203095      CURRENT  11061559    16-JAN-18
7        7       PDB_251   913203095      ORPHAN   11061957    16-JAN-18
10       10      PDB_251   913203095      ORPHAN   11062444    16-JAN-18
11       11      PDB_251   913203095      ORPHAN   11064232    16-JAN-18
12       12      PDB_251   913203095      ORPHAN   11064233    16-JAN-18

```

Clone operation will fail if multipath is disabled on the plug-in host

Description

When you perform a clone operation on the primary storage, the operation fails with an error message.

Error message

Unable to complete the build host stack operation for '/mnt/sanext4_1', reason: 'Mounting the filesystem '/mnt/sanext4_1_DBSAN50' failed, reason: mount: wrong fs type, bad option, bad superblock on /dev/sdbb1, missing codepage or helper program, or other error

Corrective action

You must configure the multipath stack in the plug-in host.

Cloning operation will fail in SAN environments on OL 7 or later or RHEL 7 or later

Description

Cloning operation will fail in storage area network (SAN) environments on Oracle Linux 7 or later or Red Hat Enterprise Linux (RHEL) 7 or later because lvm2-lvmetad service is enabled by default.

Error message

Job Failed: Failed on 'SNAPCENTER-01': Activity 'Application Clone' failed with error: CloneActivity failed PL-SCO-30000: Cloning of database with SID *SID_value* failed with error: PL-SCO-30015: Failed to get parameters from the trace file /mnt/orastadata_*SID_value*/oradata/rrdb/*debug_file*.trc with error: /mnt/orastadata_*SID_value*/oradata/rrdb/*debug_file*.trc

Corrective action

You must set the value of `use_lvmetad = 0` in `/etc/lvm/lvm.conf` and stop the lvm2-lvmetad service. Then, retry the clone operation.

Clone operation will fail due to inaccessible virtual device

Description

One of the possible reasons a clone operation can fail is because a virtual device that is listed in vCenter is inaccessible because of a permanent device loss.

For example, if the LUN ID of a newly created LUN for a clone operation matches a LUN ID that is being used by an inaccessible LUN in vCenter, the clone operation fails.

Error message

Failed to clone resources, error is Failed to Mount/Attach Disk

Corrective action

1. Using the vCenter interface, remove the inaccessible LUN.
Refer to the VMware documentation for information.
2. Using the SnapCenter interface, run the clone operation again.

Clone operation might fail or take longer time to complete with default TCP_TIMEOUT value

Description

Clone operations take longer than 30 minutes to complete when scheduled backups are running with the default TCP_TIMEOUT set.

Backup operations fails with MySQL connection error because of the delay in the TCP_TIMEOUT.

Corrective action

In the Windows registry, complete the following steps:

1. Using Registry Editor (regedit.exe), open the Windows registry.
2. Locate the following key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters
```

3. To change the available ephemeral ports from 5000 to 65534, in the Edit menu, click **New > DWORD (32-bit) Value**.
4. Add the following registry information:
 - Value Name:
MaxUserPort
 - Value:
65534
5. To change the TcpTimedWaitDelay from the default of 4 minutes to 30 seconds, in the Edit menu, click **New > DWORD (32-bit) Value**.

6. Add the following registry information:
 - Value Name:
`TcpTimedWaitDelay`
 - Value:
`30`
7. Exit the registry by closing Registry Editor.
8. Reboot the SnapCenter host.

Clone operation might fail if you are using Oracle databases 11.2.0.3 or later

Description

If you are using Oracle database 11.2.0.3 or later and the database ID for the auxiliary instance is changed using an NID script, the clone operation might fail.

Error message

```
NID-00106: LOGIN to target database failed with Oracle error ORA-01017:  
invalid username/password; logon denied
```

Corrective action

You must install the 13366202 Oracle patch.

Fails to clone an Oracle database on a volume

Description

If you do not provide a folder name, cloning an Oracle database on a volume fails. The issue occurs because the volume by default has `.snapshot` folder.

Corrective action

You should delete the `.snapshot` folder from the volume.

Recovery of a cloned database fails

Description

When you try to recover a cloned database as part of the clone operation, the operation fails. This issue occurs if the current incarnation of the database is reset to a newly detected incarnation.

The incarnation is reset if you have configured Fast Recovery Area (FRA), and if any of the auto backup of control files of the source database exist in the FRA.

Error message

```
ORA-19909: datafile 1 belongs to an orphan incarnation
```


Corrective action

You must perform the following steps:

1. Disable the auto backup of control files or ensure that the auto backup of control files does not exist in the FRA.
2. Create a backup.
3. Perform cloning using the new backup.

Recovery operation fails if the SCN specified is inconsistent

Description

During restore or clone operations, if you provide an inconsistent SCN, the recovery operation fails.

Error message

ORA-01547: warning: RECOVER succeeded but OPEN RESETLOGS would get error below
ORA-01195: online backup of file 1 needs more recovery to be consistent.

Corrective action

You should manually enter the correct SCN number.

Clone split operation stops responding

Description

The clone split operation stops responding if the SMCore service restarts.

Corrective action

Use the `Stop-SmJob` cmdlet to stop the clone split operation, and then perform the clone split operation again.

Clone split estimation fails when the aggregate does not have space

Description

If the aggregate has no space, clone split estimation fails.

Corrective action

Increase the space on the aggregate.

Clone split start operation fails

Description

The clone split start operation fails when one of the following conditions are met:

- The volume or aggregate is offline.
- A clone split operation, backup operation, or restore operation is in progress on the clone resource.
- A clone split operation is in progress on an intermediate clone.

Corrective action

- If the volume or aggregate is offline, check the status of the volume or aggregate on the storage system and bring it back online.
- If a clone split operation, backup operation, or restore operation is in progress on the clone resource, restart the clone split operation after the clone split operation, backup operation, or restore operation is complete.
- If a clone split operation is in progress on an intermediate clone, split the child clones, and then restart the clone split start operation on this clone.

Databases on which the clone split operation was performed are listed as clones

Description

During the clone split operation, if the SMCore service restarts, the status of the operation will not be sent to the server. Therefore, the databases on which the clone split operation was performed are listed as clones in the Resources page.

Corrective action

Run the split clone cmdlet again to clean the clone metadata in SnapCenter repository.

File system is not deleted during the clone delete operation

Description

While performing the clone delete operation, sometimes the file systems are not deleted.

Error message

```
NFS mount point is busy
```

Corrective action

You must increase the value of the `CLONE_DELETE_DELAY` parameter by running the following command:

```
./sccli Set-SmConfigSettings
```

Note: The `CLONE_DELETE_DELAY` parameter specifies the number of seconds to wait after completing the deletion of application clone and before starting the deletion of file system.

After modifying the value of the parameter, restart the SnapCenter Plug-in Loader (SPL) service.

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Backup and clone operations fail if stale entries of the cloned disk group exists

Description

When you perform a backup or clone operation, the operation might fail if stale entries of the cloned disk group exist in the `asm_diskgroups` parameter.

Error message

ORA-15130: diskgroup "DISKGROUP_SCO_ID" is being dismounted

Corrective action

You must clean up the `asm_diskgroups` string to remove the stale entry for `DISKGROUP_SCO_ID`.

Operations fail when there is insufficient space to create Snapshot copies

Description

ONTAP reserves space for creating Snapshot copies on volumes. If the space reserved for creating Snapshot copies is full, then Snapshot copies are not created and the operation fails.

Corrective action

You must increase the space reserved for Snapshot copies on the volumes and retry the operation.

Operations are not executed due to insufficient space in the root file system

Description

Operations might not be executed when there is insufficient space in the root file system to create logs and temporary files.

Error message

Plugin cannot accept any more jobs at this time. Job will be queued, and retried after 5 minutes.

Corrective action

You must ensure that there is sufficient space in the root file system to create logs and temporary files.

Data protection operation fails if operational lock file is not deleted

Description

While performing an operation on the database, an operational lock file (`sm_lock_dbsid`) is created in `$ORACLE_HOME/dbs` to avoid multiple operations being executed on the database. This operational lock file is automatically deleted soon after the operation is completed. However, sometimes the operational lock file might not get deleted and the next operation fails.

The cataloging operation fails if the database name is missing from the `\etc\oratab` file. The operational lock file that was created is not deleted because SID of the database cannot be retrieved.

Error message

Operation failed. The database SID `sid_value` might be in use by another SnapCenter Plug-in for Oracle Database operation.

Failed to find entry for SID `sid_value` or database name `database_name` in `/etc/oratab` file on host.

Corrective action

You must manually delete the operational lock file by performing the following steps:

1. From the command prompt, navigate to `$ORACLE_HOME/dbs`.
2. Enter the following command:

```
rm -rf .sm_lock_dbsid.
```

Data protection operation fails because of application firewall

Description

When you try to perform any data protection operation on a database, the operation might fail because of an application firewall, for example `f5`.

Error message

```
ERROR SMCore_197 PID=[2548] TID=[135] Exception in method: InvokeXML.  
System.Net.WebException: The underlying connection was closed: An  
unexpected error occurred on a receive.
```

Corrective action

You should set the time out value of the application firewall to 3 hours or more.

Operations that require backup to be mounted might fail

Description

The aggregate that is being used by the operation you are trying to perform must be assigned to the storage virtual machine (SVM) used by the database. If the aggregate is not assigned to the SVM, the operation might fail.

Error message

Failed to mount storage resource.

Corrective action

You must assign the aggregate to the SVM used by the database.

Messages in the log file display incorrect time zone

Description

An incorrect time zone is displayed in the log messages for certain versions of Java when the local time zone is not set properly in the TZ environment variable or in the Zone parameter located at `/etc/sysconfig/clock`.

Corrective action

You must perform one of the following:

- You must ensure that the correct time zone is assigned to the TZ environment variable.
For example, `TZ=America/New_York`
- If the TZ environment variable is empty, then you must ensure that the Zone parameter located at `/etc/sysconfig/clock` is set to a correct time zone.
For example, `ZONE=America/Los_Angeles`

You can also resolve this issue by changing the value of `JAVA_HOME` to JDK 7 (b72) or later.

Operations fail with command execution timeout error

Description

SnapCenter Plug-ins for Linux execute the UNIX commands to manage the file systems, Logical Volume Manager (LVM), and multipath environment. This operation sometimes takes time to complete and the operation times out.

Error Message

command execution timed out

Corrective action

You must increase the value of the `COMMAND_EXECUTION_TIMEOUT` parameter to 86400000 ms by running the following command:

```
./sccli Set-SmConfigSettings
```

Note: The `COMMAND_EXECUTION_TIMEOUT` parameter specifies the number of seconds to wait for an operation to complete.

After modifying the value of the parameter, restart the SnapCenter Plug-in Loader (SPL) service.

Related information

[*SnapCenter Software 4.1 Linux Command Reference Guide*](#)

Data protection operation fails in a non-multipath environment in RHEL 7 and later

Description

When you perform any data protection operations in a non-multipath environment in RHEL 7 and later, the operations fail with an error message.

Error message

Failed to deport the underlying stack of the file system `mount_path` as the file system belongs to volume group `volume_group_name`, and one or more physical volumes of the same volume group could not be successfully deported.

Corrective action

1. Disable or stop the logical volume manager (LVM) metadata service: `systemctl lvm2-lvmetad.service stop`
2. Change the configuration value of `use_lvmetad` from 1 to 0 in the `lvm.conf` file.
The file is located at: `/etc/lvm/` directory.
3. Restart the LVM metadata service.

Managing policies

You can create, copy, modify, view, and delete backup policies.

About this task

You can perform the following tasks related to policies:

- Create a policy.
- Modify a policy.

Note: You can change the schedule type for a policy only after you detach the policy. To change the schedule you must modify the resource group.
- Copy a policy by accepting the default name or typing a new name.
- Detach a policy from a resource group.
- Delete a policy.

Related tasks

[Creating backup policies for Oracle databases](#) on page 17

Detaching policies

You can detach policies from a resource or resource group any time that you no longer want those policies to govern data protection for the resources. You must detach a policy before you can delete it or before you modify the schedule type.

About this task

Attention: You cannot detach a policy from a resource or resource group if it is the only policy attached.

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 **Resource Group** from the **View** list.
3. Select the resource group, and then click **Modify Resource Group**.
4. In the **Policies** page of the **Modify Resource Group** wizard, from the drop-down list, clear the check mark next to the policies you want to detach.

Note: You cannot detach the last remaining policy from an individual resource because every resource must have at least one policy attached. Therefore, if you want to delete or modify the only policy attached to a resource, you must perform the following:

- a. Attach a second placeholder policy.
 - b. Detach the original policy from the resource.
5. Make any additional modifications to the resource group in the rest of the wizard, and then click **Finish**.

Modifying policies

You can modify the replication options, Snapshot copy retention settings, error retry count, or scripts information while a policy is attached to a resource or resource group. You can modify the schedule type (frequency) only after you detach a policy.

About this task

Modifying the schedule type in a policy requires additional steps because the SnapCenter Server registers the schedule type only at the time the policy is attached to a resource or resource group.

If you want to...	Then...
Add an additional schedule type	<p>Create a new policy and attach it to the necessary resources or resource groups.</p> <p>For example, if a resource group policy specifies only hourly backups and you want to add daily backups also, you can create a policy with a daily schedule type and add it to the resource group. The resource group would then have two policies: hourly and daily.</p>
Remove or change a schedule type	<ol style="list-style-type: none"> 1. Detach the policy from every resource and resource group that uses that policy. 2. Modify the schedule type. 3. Attach the policy again to all the resources and resource groups. <p>For example, if a policy specifies hourly backups and you want to change that to daily backups, you must detach the policy first.</p> <p>Note: You cannot detach the last remaining policy from a individual resource because every resource must have at least one policy attached. Therefore, if you want to modify the schedule type of the only policy attached to a resource, you must perform the following:</p> <ol style="list-style-type: none"> 1. Attach a second placeholder policy. 2. Detach the original policy from every resource and resource group that uses that policy. 3. Modify the schedule type. 4. Attach the modified policy again to all the resources and resource groups. 5. Detach the placeholder policy.

Steps

1. In the left navigation pane, click **Settings**.
2. In the **Settings** page, click **Policies**.
3. Select the policy, and then click **Modify**.

4. Modify the information, and then click **Finish**.

Deleting policies

If you no longer require policies, you might want to delete them.

Before you begin

You must have detached the policy from resource groups if the policy is associated with any resource group.

Steps

1. In the left navigation pane, click **Settings**.
2. In the **Settings** page, click **Policies**.
3. Select the policy, and then click **Delete**.
4. Click **Yes**.

Managing resource groups

You can create, modify, and delete resource groups. You can also perform backup and verification operations on resource groups.

About this task

You can perform the following tasks related to resource groups:

- Create a resource group.
- Modify a resource group by selecting the resource group and clicking **Modify Resource Group** to edit the information you provided while creating the resource group.

Note: You can change the schedule while modifying the resource group. However, to change the schedule type you must modify the policy.

Note: If you remove resources from a resource group, the backup retention settings defined in the policies currently attached to the resource group will continue to be applied to the removed resources.

- Create a backup of a resource group.
- Create a clone of a backup.
You can clone from the existing backups of SQL, Oracle, Windows file systems, custom applications, and SAP HANA database resources or resource groups.
- Create a clone of a resource group.
This operation is supported only for SQL resource groups (which contains only databases). You can configure a schedule for cloning a resource group (clone lifecycle).
- Prevent scheduled operations on resource groups from starting.
- Delete a resource group.

Related tasks

[Creating resource groups and attaching policies for Oracle databases](#) on page 21

[Backing up Oracle database resource groups](#) on page 25

Stopping and resuming operations on resource groups

You can temporarily disable scheduled operations from starting on a resource group. Later when you want, you can enable those operations.

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 **Resource Group** from the **View** list.
3. Select the resource group and click **Maintenance**.
4. Click **OK**.

After you finish

If you want to resume operations on the resource group that you had put on maintenance mode, select the resource group and click **Production**.

Deleting resource groups

You can delete a resource group if you no longer need to protect the resources in the resource group. You must ensure that resource groups are deleted before you remove plug-ins from SnapCenter.

Before you begin

If you are managing SQL or Windows file system resources, you must have manually deleted all clones created for any of the resources in the resource group.

About this task

You can optionally force the deletion of all backups, metadata, policies, and Snapshot copies associated with the resource group.

Note: In a SnapVault relationship, the last Snapshot copy cannot be deleted; therefore, the resource group cannot be deleted. Before deleting a resource group that is part of a SnapVault relationship, you must remove the SnapVault relationship and then delete the last Snapshot copy.

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 **Resource Group** from the **View** list.
3. Select the resource group, and then click **Delete**.
4. Optional: Select the **Delete backups and detach policies associated with this Resource Group** check box to remove all backups, metadata, policies, and Snapshot copies associated with the resource group.
5. Click **OK**.

Managing backups

You can rename and delete backups. You can also delete multiple backups simultaneously.

Renaming backups

You can rename backups if you want to provide a better name to improve searchability.

Steps

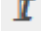
1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page either select the resource or resource group from the **View** drop-down list.
3. Select the resource or resource group from the list.

The resource or resource group topology page is displayed. If the resource or resource group is not configured for data protection, the Protect wizard is displayed instead of the topology page.

4. From the **Manage Copies** view, select **Backups** from the primary storage systems.

You cannot rename the backups that are on the secondary storage system.

If you are using SnapCenter Plug-ins Package for Linux and have cataloged the backup using Oracle Recovery Manager (RMAN), you cannot rename those cataloged backups.

5. Select the backup, and then click .
6. In the **Rename backup as** field, enter a new name and click **OK**.

Deleting backups

You can delete backups if you no longer require the backup for other data protection operations.

Before you begin

You must have deleted the associated clones before deleting a backup.

If a backup is associated with a cloned resource, you cannot delete the backup.


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 or resource group from the list.

The resource or resource group topology page is displayed.

4. From the **Manage Copies** view, select **Backups** from the primary storage systems.

You cannot delete the backups that are on the secondary storage system.

5. Select the backup, and then click .
6. Click **OK**.

Note: If you have some stale database backups in SnapCenter which do not have corresponding backups on the storage system, you must use `remove-smbbackup` command to clean up these stale backup entries. If the stale backups were cataloged, they will be uncataloged from the recovery catalog database.

Related information

[SnapCenter Software 4.1 Linux Command Reference Guide](#)

Managing clones

You can view and delete clones.

Deleting clones

You can delete clones if you find them no longer necessary.


About this task

A clone that has been cloned again cannot be deleted. For example, if the production database *db1* is cloned to *db1_clone1* and subsequently cloned to *db1_clone2*, and you decide that you want to delete *db1_clone1*, you must first delete *db1_clone2*, and then delete *db1_clone1*.

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 or resource group from the list.

The resource or the resource group topology page is displayed.

4. From the **Manage Copies** view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
5. Select the clone, and then click .
6. Click **OK**.

Copyright

Copyright © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277