



SnapCenter® Software 1.1

Data Protection Guide

For Oracle Databases

May 2016 | 215-10972_CO
doccomments@netapp.com

 **NetApp®**

Contents

Data protection using SnapCenter Plug-in for Oracle Database	6
What you can do with the SnapCenter Plug-in for Oracle Database	6
SnapCenter Plug-in for Oracle Database features	6
SnapCenter components	7
How resources, datasets, and policies are used in data protection	10
Data protection workflow for Oracle databases	12
Preparing for data protection	13
Prerequisites for using the SnapCenter Plug-in for Oracle Database	13
Storage types supported by SnapCenter Plug-in for Oracle Database	14
Logging in to SnapCenter	15
Backing up Oracle databases	16
Defining a backup strategy for Oracle databases	17
Supported Oracle database configurations for backups	17
Types of backup supported	18
Why to specify preferred nodes in RAC setup	19
When to back up your Oracle databases	19
Which backup naming convention to use	19
Backup retention options	20
When to schedule verification jobs	20
How SnapCenter works with SnapMirror and SnapVault technologies	20
Backup copy verification using the primary or secondary storage volume	21
Supported prescripts and postscripts	21
Determining whether Oracle databases are available for backup	21
Creating verification policies for Oracle databases	22
Creating backup policies for Oracle databases	23
Creating backup datasets and attaching policies for Oracle databases	27
Backing up resources	29
Backing up resources on demand	29
Backing up datasets on demand	30
Monitoring backup operations from the Jobs page	30
Verifying Oracle database backups	32
Creating verification policies for Oracle databases	32
Verifying an Oracle database backup	33
Monitoring verification operations from the Jobs page	34
Mounting and unmounting database backups	36
Mounting a database backup	36
Unmounting a database backup	37
Restoring and recovering Oracle databases	38
Defining a restore and recovery strategy for Oracle databases	38
Backups supported for restore and recovery operations	39

Limitations related to restore and recovery operations	40
Types of restore methods supported for Oracle databases	40
Types of restore operations supported for Oracle databases	40
Types of recovery operations supported for Oracle databases	41
Sources and destinations for a restore operation	41
Restoring and recovering an Oracle database	41
Monitoring restore operations from the Jobs page	44
Cloning Oracle database backups	46
Defining a clone strategy for Oracle databases	46
Backups supported for cloning	47
Type of cloning supported	47
Cloning an Oracle database backup	47
Monitoring clone operations from the Jobs page	51
Troubleshooting data protection operations	52
Backup operation fails when initiated before starting the database instance	52
Backup fails with license issue	52
Backup fails during the file system discovery on a VM	52
Backup operation fails during the storage discovery process	53
Registering backup activity fails during the backup operation	53
Verification of datafiles backup fails	54
Backup verification fails	54
Disk paths are not included in the asm_diskstring database parameter	54
Unable to change the database state from shutdown to mount	55
Restore operation of datafiles and control files fail	55
Restore from a secondary SnapMirror or SnapVault volume fails	55
Cloning operation on the primary storage fails	56
Cloning operation fails in SAN environments in OL 7 or later or RHEL 7 or later	56
Recovery of a cloned database fails	56
File system is not deleted during the clone delete operation	57
Backup and clone operations fail if stale entries of the cloned disk group exists	57
Operations fail when there is insufficient space to create Snapshot copies	58
Operations are not executed due to insufficient space in the root file system	58
Data protection operation fails if operational lock file is not deleted	58
Messages in the log file display incorrect time zone	59
Operations fail with command execution timeout error	59
Data protection operation fails in a non-multipath environment in RHEL 7 and later	59
Managing Oracle database datasets	61
Modifying datasets	61
Stopping operations on datasets temporarily	61
Resuming operations on datasets	62
Deleting datasets for Oracle databases	62
Managing Oracle database policies	63
Modifying policies	63

Copying policies	63
Viewing policy details	64
Detaching policies from a dataset	64
Deleting policies	64
Managing backups	65
Renaming or deleting backup copies	65
Managing clones	66
Viewing clone details	66
Deleting clones	66
Backing up, restoring, and cloning using Linux commands	67
Backing up Oracle databases using Linux commands	67
Restoring and recovering Oracle databases using Linux commands	68
Cloning Oracle database backups using Linux commands	69
Where to go next	70
Copyright information	71
Trademark information	72
How to send comments about documentation and receive update notifications	73
Index	74

Data protection using SnapCenter Plug-in for Oracle Database

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 automates the backup, verification, mounting, unmounting, restore, recovery, and cloning of Oracle databases in your SnapCenter environment.

The Plug-in for Oracle requires the SnapCenter Plug-in for UNIX to perform backup, verification, restore, recovery, and cloning of Oracle databases that are running on Linux hosts.

When the Plug-in for Oracle is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance.

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

Related information

[SnapCenter Software 1.1 Installation and Setup Guide](#)

[SnapCenter Software 1.1 Getting Started Guide](#)

What you can do with the SnapCenter Plug-in for Oracle Database

You can use the Plug-in for Oracle to back up, verify, restore, recover, mount, unmount, and clone Oracle databases and their resources.

- Back up datafiles, control files, and archive log files.
- Restore database resources, including databases, tablespaces, container databases (CDBs), and pluggable databases (PDBs) and recover them up to a point-in-time.
- Create clones of production databases up to a point-in-time.
- Verify backups immediately or defer the verification.
- Mount and unmount log backups for recovery operation.
- Schedule backup, verification, and clone operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, verification, restore, and clone operations.

SnapCenter Plug-in for Oracle Database features

The Plug-in for Oracle integrates with the Oracle database on the Linux host and with NetApp technologies on the storage system.

Unified graphical user interface

The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, recovery, and clone operations across plug-ins, use centralized reporting, use at-a-

glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

Automated central administration

You can schedule backup and clone operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment with periodic email alerts.

Nondisruptive NetApp Snapshot copy technology

SnapCenter uses NetApp Snapshot copy technology with the Plug-in for Oracle and Plug-in for UNIX to back up databases. Snapshot copies consume minimal storage space.

The Plug-in for Oracle also offers the following benefits:

- Support for backup, restore, clone, mount, unmount, and verification workflows
- Automatically discover Oracle databases configured on the host
- RBAC-supported security and centralized role delegation
You can also set the “Run As” credentials so that the authorized SnapCenter users have application-level permissions.
- Creation of space-efficient and point-in-time copies of production databases for testing or data extraction by using NetApp FlexClone technology
A FlexClone license is required on the storage system where you want create the clone.
- Support for consistency group (CG) feature of Data ONTAP as part of creating backups in SAN and ASM environments
- Nondisruptive and automated backup verification
- Capability to run multiple backups simultaneously across multiple database hosts
In a single operation, Snapshot copies are consolidated when databases in a single host share the same volume.
- Support for physical and virtualized infrastructures
- Support for NFS, iSCSI, Fibre Channel (FC), RDM, VMDK over NFS and VMFS, and ASM over NFS and SAN
- Support for the Selective LUN Map (SLM) feature of Data ONTAP
Enabled by default, the SLM feature periodically discovers the LUNs that do not have optimized paths and fixes them. You configure SLM by modifying the parameters in the `scu.properties` file located at `/var/opt/snapcenter/scu/etc`.
 - You can disable this by setting the value of `ENABLE_LUNPATH_MONITORING` to **false**.
 - You can specify the frequency in which the LUN paths will be fixed automatically by assigning the value (in hours) to `LUNPATH_MONITORING_INTERVAL`.

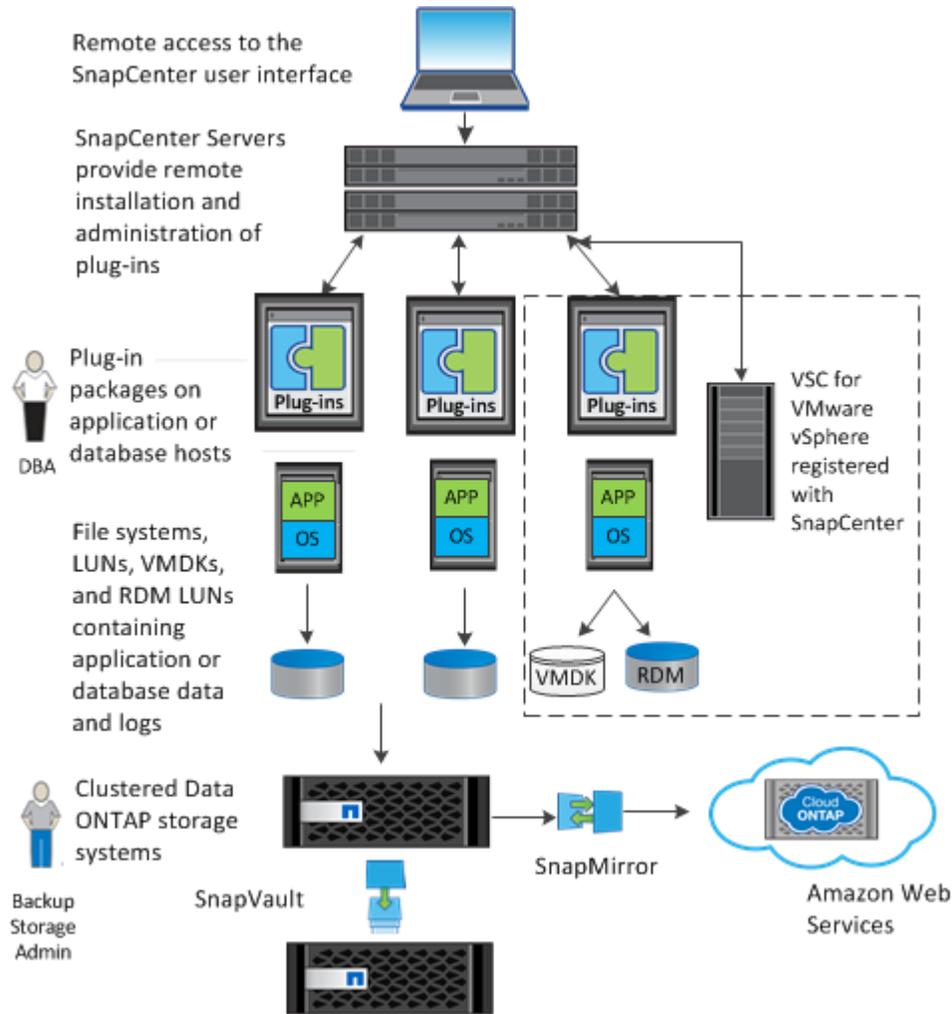
For information about SLM, see the Data ONTAP administration information.

[Clustered Data ONTAP 8.3 SAN Administration Guide](#)

SnapCenter components

SnapCenter consists of the SnapCenter Server, the SnapCenter Plug-ins Package for Windows, and the SnapCenter Plug-ins Package for Linux. Each of these packages contain plug-ins to SnapCenter. SnapCenter also interacts with Virtual Storage Console for VMware vSphere (VSC) to provide support for database backup, restore, recovery, and cloning on raw device mappings (RDMs) and

virtual machine disks (VMDKs). When you are installing and configuring SnapCenter, it is helpful to understand its components.



SnapCenter Server

The SnapCenter Server includes a web server, a centralized HTML5-based user interface, PowerShell cmdlets, APIs, and the SnapCenter repository. SnapCenter enables load balancing, high availability, and horizontal scaling across multiple SnapCenter Servers within a single user interface. You can accomplish high availability by using Network Load Balancing (NLB) and Application Request Routing (ARR) with SnapCenter. For larger environments with thousands of hosts, adding multiple SnapCenter Servers can help balance the load.

The SnapCenter platform is based on a multitiered architecture that includes a centralized management server (SnapCenter Server) and a SnapCenter host agent (SMCore):

- If you are using the SnapCenter Plug-ins Package for Windows, the SMCore runs on the SnapCenter Server and Windows plug-in host. The SnapCenter Server communicates with the Windows plug-ins through the SMCore.
- If you are using the SnapCenter Plug-ins Package for Linux, the SMCore running on the SnapCenter Server host communicates with the SnapCenter Plug-in Loader (SPL) running on the Linux host to perform different data protection operations.

SnapCenter enables centralized application resource management and easy data protection job execution through the use of datasets and policy management (including scheduling and retention settings). SnapCenter provides unified reporting through the use of a Dashboard, multiple reporting options, job monitoring, and log and event viewers.

Information about SnapCenter operations is stored in the SnapCenter repository.

SnapCenter Plug-ins Package for Windows

This installation package includes the following plug-ins:

SnapCenter Plug-in for Microsoft SQL Server

The Plug-in for SQL Server is a host-side component of the NetApp storage solution offering application-aware backup management of Microsoft SQL Server databases. With the plug-in installed on your SQL Server host, SnapCenter automates Microsoft SQL Server database backup, restore, and clone operations.

If you want to perform data protection operations on SQL Servers that are on VMDKs or RDMs, you must register VSC with SnapCenter using the SnapCenter Add Hosts wizard or the VSC Configure SnapCenter Server dialog box.

SnapCenter Plug-in for Microsoft Windows

The Plug-in for Windows provides storage provisioning, Snapshot copy consistency, and space reclamation for Windows hosts. With the plug-in installed on your Windows host, you can use SnapCenter to create and resize disks, initiate iSCSI sessions, manage igroups, and manage SMB shares. The Plug-in for Windows is a required component of the Plug-in for SQL Server workflows.

Support is provided for provisioning SMB shares only. You cannot use SnapCenter to back up SQL Server databases on SMB shares.

Some features are not supported on all operating systems. See information about host requirements in the installation prerequisites.

SnapCenter Plug-ins Package for Linux

This installation package includes the following plug-ins:

SnapCenter Plug-in for Oracle Database

The Plug-in for Oracle is a host-side component of the NetApp integrated storage solution offering application-aware backup management of Oracle databases. With the Plug-in for Oracle installed on your Oracle host, SnapCenter automates backup, restore, recovery, verify, mount, unmount, and clone operations.

If you want to perform data protection operations on Oracle databases that are on VMDKs or RDMs, you must register VSC with SnapCenter using the SnapCenter Add Hosts wizard or the VSC Configure SnapCenter Server dialog box.

Note: You can use the Plug-in for Oracle to manage Oracle databases for SAP as well. However, SAP BR*Tools integration is not supported.

SnapCenter Plug-in for UNIX

The Plug-in for UNIX handles the underlying host storage stack and enables you to perform backup, restore, clone, mount, and unmount operations on Oracle databases that are running on a Linux host by working in conjunction with the Plug-in for Oracle. The Plug-in for UNIX supports the Network File System (NFS) and storage area network (SAN) protocols on a storage system that is running clustered Data ONTAP.

Virtual Storage Console for VMware vSphere

Virtual Storage Console for VMware vSphere (VSC) is a vCenter Server plug-in that provides end-to-end lifecycle management for virtual machines in VMware environments using NetApp storage systems. Integrating SnapCenter with VSC enables SnapCenter to perform backup, clone, and restore operations on database environments (SQL Server or Oracle database) that are backed by VMware datastores, VMDKs, or RDMs.

VSC uses SnapCenter to perform backup, restore, and clone operations for storage systems running clustered Data ONTAP 8.2.2 or later. Integrating VSC with SnapCenter enables VSC backup and recovery to use the SnapCenter database to store metadata. This greatly increases the scalability of VSC. This integration also enables the use of backup policies, enabling you to create and reuse policies in VSC, even across multiple VSC environments.

You must register VSC with SnapCenter using the SnapCenter Add Hosts wizard or the VSC Configure SnapCenter Server dialog box.

If you used VSC to create policies and datasets and registered VSC with SnapCenter, you can view these policies and datasets in SnapCenter. You can also view VSC-based backup, clone, and restore operations in the SnapCenter Jobs view.

VSC is not included in the SnapCenter installation.

Note: You need not register VSC with SnapCenter if you are using Plug-in for SQL Server (and your SQL Server environment uses an iSCSI initiator) or Plug-in for Oracle (and your Oracle environment uses NFS or an iSCSI initiator). If you do not register VSC with SnapCenter, a message appears in the SnapCenter Hosts view that indicates you should configure the hypervisor. For these conditions, you do not need to configure the hypervisor.

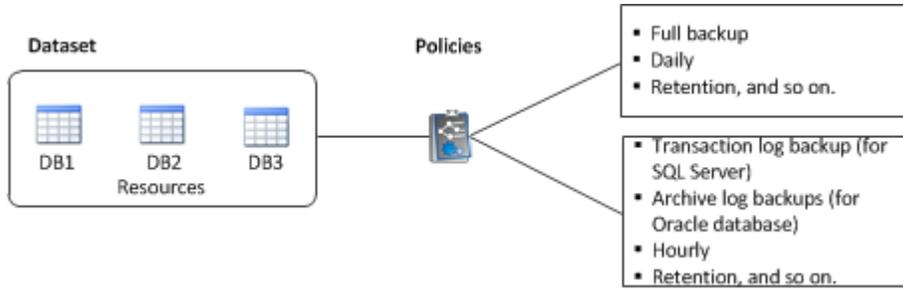
How resources, datasets, and policies are used in data protection

Before you use SnapCenter, it is helpful to understand basic concepts related to the backup, clone, and restore operations you want to perform. You interact with resources, datasets, and policies for different operations.

- *Resources* are typically databases that you back up or clone.
However, depending on your environment, resources might be database instances, Microsoft SQL server availability groups, Oracle databases, Oracle RAC databases, or anything that you want to back up or clone.
- A SnapCenter *dataset* is a collection of resources on a host or cluster.
When you perform an operation on a dataset, you perform that operation on the *resources* defined in the dataset.
- The *policies* specify the schedule, copy retention, replication, scripts, and other characteristics of data protection operations.
You select the policy when you perform the operation.

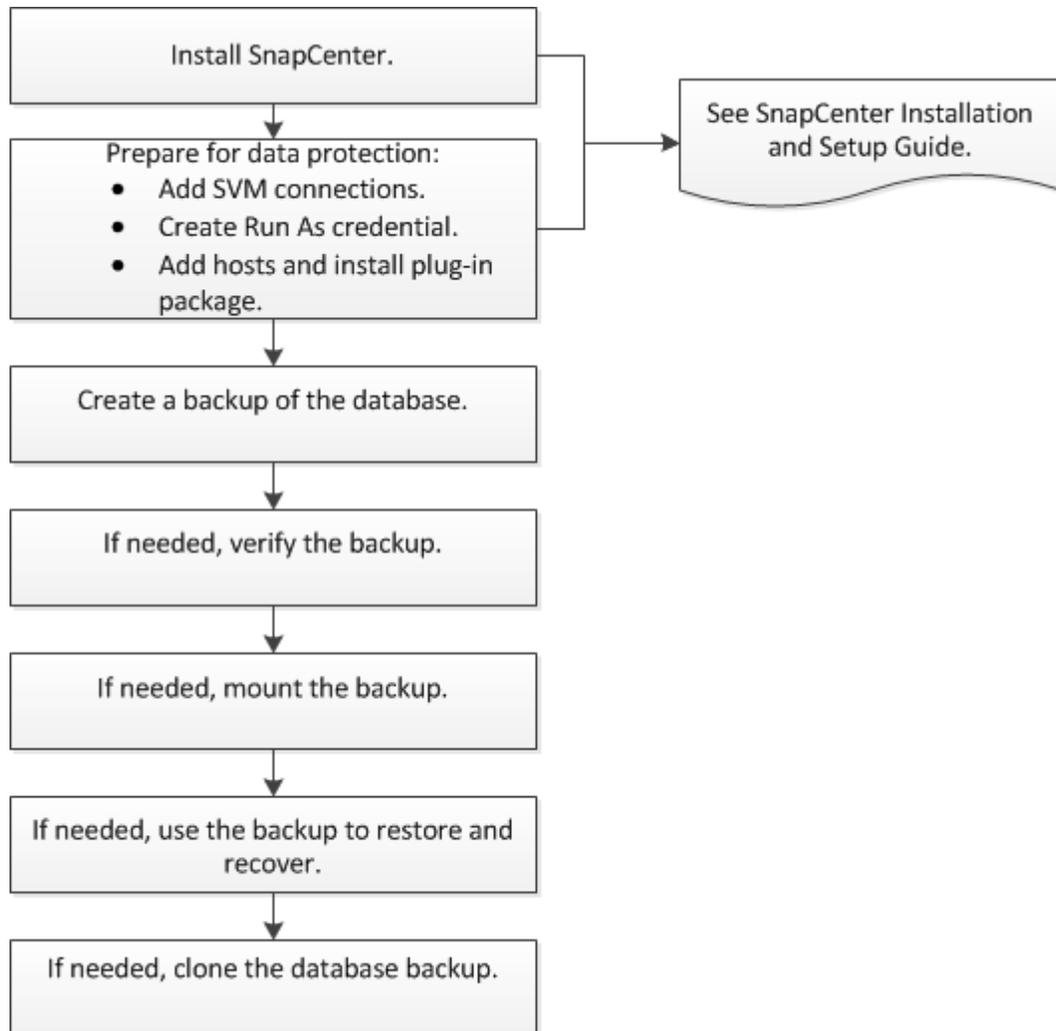
Think of a dataset as defining *what* you want to protect and a policy as defining *how* you want to protect it. If you are backing up all databases of a host, for example, you might create a dataset that includes all the databases in the host. You could then attach two policies to the dataset: one that performs a full backup daily and another that performs transaction log (for SQL) or archive log backups (for Oracle database) hourly.

The following image illustrates the relationship between resources, datasets, and policies:



Data protection workflow for Oracle databases

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



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.

Related information

[SnapCenter Software 1.1 Installation and Setup Guide](#)

Prerequisites for using the SnapCenter Plug-in for Oracle Database

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

- Install and configure SnapCenter Server.
- Log in to SnapCenter Server.
- Configure the SnapCenter environment by adding SVM connections and creating Run As credentials.
 - Note:** The Run As credential must be created for the root user or equivalent of the Linux host on which the SnapCenter Plug-ins Package for Linux will be installed.
- Add hosts, install the plug-ins, and discover the resources.
- Register Virtual Storage Console (VSC) with SnapCenter Server if you are using SnapCenter Server to protect Oracle databases that reside on VMware RDM LUNs or VMDKs. The registration of VSC enables SnapCenter Server to communicate with VMware vSphere.
- Install Java 1.7 or Java 1.8 on your Linux host.
- If you have multiple data paths (LIFs) or dNFS configuration, you can perform the following using SnapCenter CLI on the database host:
 - By default all the IP addresses of the database host is 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 path (LIF) 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 information to run the CLIs, see the [SnapCenter Software 1.1 Linux Command Reference Guide](#).

- If you are having Oracle databases on NAS or 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 are using Oracle database on SAN environments, ensure that the SAN environment is configured as per the recommendation in the [Recommended Host Settings for Linux Unified Host Utilities 7.0](#) and [Using Linux Hosts with Data ONTAP Storage](#).
- If you are using Oracle database on LVM in Oracle Linux or RHEL 6.6 or 7.0 operating systems, install the latest version of Logical Volume Management (LVM).
- If you were using SnapManager for Oracle, you can use the migration feature to create policies and datasets in SnapCenter for the SnapManager for Oracle profiles and operations performed using those profiles.
- Set up SnapMirror and SnapVault, if you want backup replication.

Related information

[SnapCenter Software 1.1 Installation and Setup Guide](#)

[SnapCenter Software 1.1 SnapManager Migration Guide](#)

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.

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

Machine	Storage type	Support notes
Physical server	FC-connected LUNs	
	iSCSI-connected LUNs	
	NFS-connected volumes	
VMware ESX	RDM LUNs connected by an FC or iSCSI HBA	You must register Virtual Storage Console (VSC) with SnapCenter before you use SnapCenter to back up databases on RDM LUNs.
	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	
	VMDKs on VMFS or NFS datastores	You must register VSC with SnapCenter before you use SnapCenter to back up databases on VMDKs.
	NFS volumes connected directly to the guest system	

Logging in to SnapCenter

Through SnapCenter role-based access control, users or groups are assigned roles and resources. When you log in to the SnapCenter graphical user interface, you log in with an Active Directory account.

About this task

The SnapCenter graphical user interface (GUI) URL is configured based on information you provide during installation. It is useful to know where to find it after you complete the SnapCenter installation.

During the installation, the SnapCenter Server Install wizard creates a shortcut and places it on your local host desktop. Additionally, at the end of the installation, the Install wizard provides the SnapCenter URL, which you can copy if you want to log in from a remote system.

The default GUI URL is a secure connection to 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`).

In addition to using the SnapCenter GUI, you can use either PowerShell cmdlets to script configuration, backup, restore, verification, and clone operations for Microsoft SQL databases or use the SnapCenter command line interface (CLI), *scccli*, to script configuration, backup, restore, verification, mount, unmount, and clone operations for Oracle database. For details, see the SnapCenter cmdlet or SnapCenter CLI documentation.

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. Choose which user you want to log in as:

If you want to ...	Do the following ...
Log in as the SnapCenter administrator	Enter the domain user with local administrator credentials provided during the SnapCenter installation.
Log in as a SnapCenter user	Enter your user credentials: <i>Domain\UserName</i>

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

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

Related information

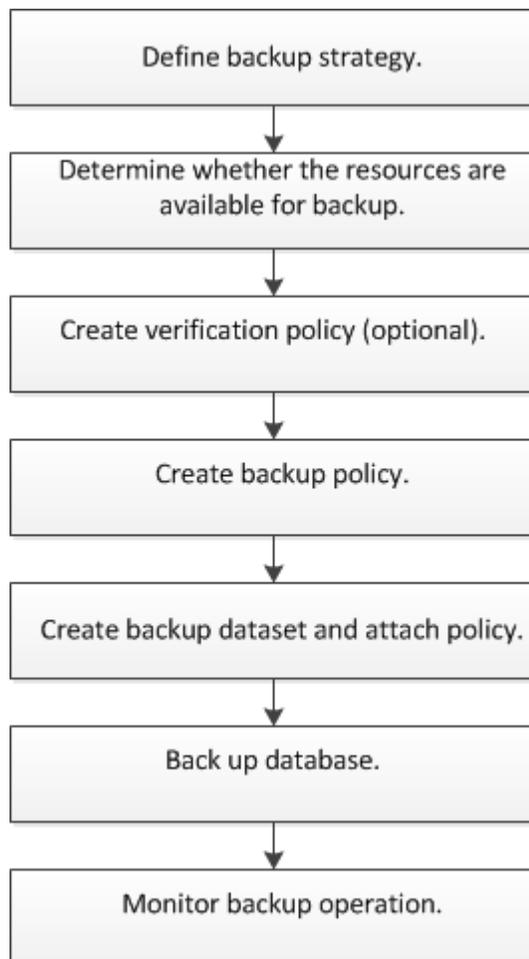
[SnapCenter Software 1.1 Linux Command Reference Guide](#)

Backing up Oracle databases

The backup workflow includes planning, identifying the resources for backup, creating verification and backup policies, creating datasets 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:



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

[SnapCenter Software 1.1 Linux Command Reference Guide](#)

Related concepts

[Backing up, restoring, and cloning using Linux commands](#) on page 67

Defining a backup strategy for Oracle databases

Defining a backup strategy before you create your backup jobs ensures that you have the backups that you require to successfully restore or clone your databases. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine the supported database configurations.
2. Decide the type of backup that you require.
3. Decide on which node in a RAC environment you want to create your backup.
4. Determine when you should back up your databases.
5. Decide how to name your backups.
6. Determine when you should verify the backup copies.
7. Decide whether you want to verify the backup copies using the primary or secondary storage volume.
8. Determine the retention period for the Snapshot copies on the source storage system and the SnapMirror destination.
9. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long term retention.
10. Determine the supported prescripts and postscripts.

Related references

Supported Oracle database configurations for backups on page 17

Types of backup supported on page 18

Why to specify preferred nodes in RAC setup on page 19

When to back up your Oracle databases on page 19

Which backup naming convention to use on page 19

When to schedule verification jobs on page 20

Backup retention options on page 20

How SnapCenter works with SnapMirror and SnapVault technologies on page 20

Supported prescripts and postscripts on page 21

Supported Oracle database configurations for backups

SnapCenter supports backup of different Oracle database configurations.

- Oracle 11g Standalone

- Oracle 11g Data Guard
- Oracle 11g Active Data Guard
- Oracle 11g Real Application Clusters (RAC)
- Oracle 12c Standalone Legacy
- Oracle 12c Standalone Container Database (CDB)
- Oracle 12c Data Guard
- Oracle 12c Active Data Guard
- Oracle 12c RAC
- Oracle database on Automatic Storage Management (ASM)
You can perform backup, restore, and clone operations on Oracle databases on ASM, with or without ASMLib. However, ASM on Raw device mapping (RDM) and virtual machine disk (VMDK) is not supported.

Note: Before creating a backup of Data Guard or Active Data Guard database, the managed recovery process (MRP) is stopped and once the backup is created, MRP is started.

Related information

[NetApp Interoperability](#)

Types of backup supported

Backup type specifies the type of backup that you want to create. SnapCenter supports online and offline backup types for Oracle databases.

Online backup

A backup that is created when the database is in the online state is called an *online backup*. Also called a *hot backup*, an online backup enables you to create a backup of the database without shutting it down.

As part of online backup, you can create a backup of the following files:

- Datafiles and control files only
- Archive log files only (the database is not brought to backup mode in this scenario)
- Full database that includes datafiles, control files, and archive log files

Note: When a backup is created, quiescing and unquiescing operations are performed for file systems of ext3, ext4, and xfs types only and not for NFS.

Offline backup

A backup created when the database is either in a mounted or shutdown state is called an *offline backup*. An offline backup is also called a *cold backup*. You can include only datafiles and control files in offline backups. You can create either an offline mount or offline shutdown backup.

- When creating an offline mount backup, you must ensure that the database is in a mounted state. If the database is in any other state, the backup operation fails.
- When creating an offline shutdown backup, the database can be in any state. The database state is changed to the required state to create a backup. After creating the backup, the database state is reverted to the original state.

Why to specify preferred nodes in RAC setup

In Oracle Real Application Clusters (RAC) setup, you can specify the preferred nodes on which the backup operation will be performed. If you do not specify the preferred node, SnapCenter automatically assigns a node as the preferred node and backup is created on that node.

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

For example, the RAC database cdbrac has three instances, namely cdbrac1 on node1, cdbrac2 on node2, and cdbrac3 on node3. The node1 and node2 are configured to be the preferred nodes with node2 being the first preference and node1 as the second preference. When you perform a backup operation, the operation will be first attempted on node2 because it is the first preferred node. If node2 is not in the state to back up which could be due to multiple reasons such as the plug-in agent is not running on the host, or the database instance on the host is not in the required state for the specified backup type; then the operation will be attempted on node1. The node3 will not be used for backup because it is not on the list of preferred nodes.

For information about how to specify the preferred nodes, see the [SnapCenter Software 1.1 Administration Guide](#).

Required database state

The RAC database instances on the preferred nodes must be in the required state for the backup to complete successfully.

- One of the RAC database instances in the configured preferred nodes must be in the open state to create an online backup.
- One of the RAC database instances in the configured preferred nodes must be in the mount state and all other instances including other preferred nodes must be in the mount state or lower to create an offline mount backup.
- RAC database instances can be in any state, but you must specify the preferred nodes to create an offline shutdown backup.

When to back up your Oracle databases

The most critical factor in determining a database backup schedule is the rate of change for the database. The more often you back up your databases, the fewer archive logs SnapCenter has to use for restoring, which can result in faster restore operations.

You might back up a heavily used database every hour, while you might back up a rarely used database once a day. Other factors include the importance of the database to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA and RPO contribute to the data protection strategy.

Which backup naming convention to use

You can either use the default Snapshot copy naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot copy names that helps you identify when the copies were created.

The Snapshot copy uses the following default naming convention:
datasetname_hostname_timestamp

You should name your backup datasets logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- `dts1` is the dataset name.
- `mach1x88` is the host name.
- `03-12-2015_23.17.26` is the date and timestamp.

Alternatively, you can specify the Snapshot copy name format, while creating datasets by selecting **Use custom name format for Snapshot copy**. For example, `customtext_dataset_policy_hostname` or `dataset_hostname`. By default, the time stamp suffix is added to the Snapshot copy name.

Backup retention options

You can choose either the number of days for which to retain backup copies or specify the number of backup copies you want to retain, up to a maximum of 255 copies. For example, your organization might require that you retain 10 days of backup copies or 130 backup copies.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

Note: For long-term retention of backup copies, you should use SnapVault backup.

When to schedule verification jobs

Although SnapCenter can verify backups immediately after it creates them, doing so can significantly increase the time required to complete the backup job and is resource intensive. Hence, it is almost always best to schedule verification in a separate job for a later time. For example, if you back up a database at 5:00 p.m. every day, you might schedule verification to occur an hour later at 6:00 p.m.

For the same reason, it is usually not necessary to run backup verification every time you perform a backup. Performing verification at regular but less frequent intervals is usually sufficient to ensure the integrity of the backup. A single verification job can verify multiple backups at the same time.

How SnapCenter works with SnapMirror and SnapVault technologies

With the SnapCenter Server, you can use SnapMirror to create mirror copies of backups on another volume. You can use SnapCenter Server with SnapVault technology to perform disk-to-disk backup replication for standard compliance.

Before you can use these technologies in SnapCenter, you must use System Manager or the ONTAP CLIs to configure a data-protection relationship between the source and destination storage volumes and initialize the relationship.

Subsequently, when you are preparing a backup, you set up a policy in which you can choose to use the SnapMirror and/or SnapVault technology.

Related information

[SnapCenter Software 1.1 Installation and Setup Guide](#)

[SnapCenter Software 1.1 Administration Guide](#)

Backup copy verification using the primary or secondary storage volume

You can verify backup copies on the primary storage volume or on either the SnapMirror or SnapVault secondary storage volume. Verification using a secondary storage volume reduces load on the primary storage volume.

When you verify on the primary or secondary storage volume, both the primary and the secondary Snapshot copies are marked as verified.

SnapRestore license is required to verify backup copies on SnapMirror and SnapVault secondary storage volume.

Supported prescripts and postscripts

You can use prescripts and postscripts as part of backup, verify, restore, mount, and clone operations. These scripts enable automation either before or after your data protection job. Before you set up your prescripts and postscripts, you should understand some of the requirements for creating these scripts.

Supported script types

You can write your script using any scripting language available in the Linux environment. You must set the executable permission for the scripts.

While writing your own script, you must specify the following return codes in your script:

- Return codes 0 and 1 implies that the script is executed successfully and the workflow continues
- Return code 2 implies that the script is executed with a warning and the workflow continues
- Return code 3, 4, or any other value implies that script is not executed successfully and the workflow fails

Script path location

You must ensure that the prescripts and postscripts should be located either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path.

Supported prescript and postscript arguments

You can only specify custom arguments for passing to the scripts. These custom arguments must be a space separated list.

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 datasets to perform data protection operations after you discover the databases that are available. Determining the available resources also verifies that the plug-in installation has been completed successfully.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, creating SVM connections, and adding “Run As” credentials.
- If the databases reside on a VMDK or RDM, you must have configured Virtual Storage Console (VSC) with the SnapCenter Server.

About this task

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

Steps

1. In the left navigation pane, click **Inventory**.
2. Click **Oracle Database**.
3. Select the host from the **Host** drop-down list.

The databases that are installed on the host are displayed along with general information about the databases.

Note: The Oracle version is not displayed if the database is in shutdown state.

4. Click **Refresh Resources**.

The newly added or deleted Oracle databases are updated to the SnapCenter Server inventory.

Note: If you have enabled Oracle database authentication while configuring the database, the discovered database is shown with a lock icon. If the lock icon appears, you must specify the database credentials for successfully adding the database to a dataset. You can specify the database credentials using **Configure Database**.

Related information

[SnapCenter Software 1.1 Administration Guide](#)

Creating verification policies for Oracle databases

You must create a verification policy to verify your backup. A verification policy is a set of rules that governs how you manage and schedule verification operations. Additionally, you can specify replication, prescript, and postscript details.

Steps

1. In the left navigation pane, click **Policies**.
2. Click **New > Oracle Policy > Verification**.
3. In the **Name** page, enter the policy name and description.
4. In the **Schedule** page, select the frequency of the verification operation from the **Schedule type** type drop-down list:

If you select...	Then...
OneTime	Specify the start date.
Hourly	Specify the start date, the interval when the operation should occur, and optionally an expiry date.
Daily	Specify the start date, the days and interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.
Weekly	Specify the start date, the interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.

If you select...	Then...
Monthly	Specify the start date, the month, days, and interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.

If you do not want to schedule the verification operation, you must select **None**.

5. In the **Replication** page, select **Verify on secondary storage** if you want to verify the backup on a secondary storage volume.

The Snapshot copies are verified on the secondary storage volume and the primary Snapshot copies are marked as verified.

You must ensure that the SVM settings of the secondary storage are included in the SnapCenter Server settings if you are verifying on the secondary storage.

If you want to verify the backups soon after the backup operation is completed, you must select the secondary replication option while creating the backup policy.

6. 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 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.

7. In the **Options** page, enter the number of most recent backups you want to verify.

Example

If three backups were created as part of the schedule and you enter 2 as the number of backups you want to verify; then of the three backups, 2 most recent backups will be verified.

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

Related concepts

[How resources, datasets, and policies are used in data protection](#) on page 10

Related information

[SnapCenter Software 1.1 Installation and Setup Guide](#)

Creating backup policies for Oracle databases

Before you use SnapCenter to back up Oracle database resources, you must create a backup policy for the dataset that you want to back up. A backup policy is a set of rules that governs how you manage, schedule, and retain backups. Additionally, you can specify the replication, script, and backup type settings. Creating a policy saves time when you want to reuse the policy on another dataset.

Before you begin

- You must have defined your backup strategy.
- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, discovering databases, and creating Storage Virtual Machine (SVM) connections.

- You must have created a verification policy if you want to verify the backup soon after the backup operation is completed.
- The SnapCenter administrator must have assigned the SVMs for both the source and destination volumes to you if you are replicating Snapshot copies to a mirror or vault.

Steps

1. In the left navigation pane, click **Policies**.
2. Click **New > Oracle Policy > Backup**.
3. In the **Name** page, enter the policy name and description.
4. In the **Backup Type** page, select the backup type:
 - 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 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.
 - If you want to **verify the backup** soon after the backup operation is completed, perform the following steps:
 - a. Select **Run verification after backup**.
 - b. Select the verification policy.
 - If you want to prune archive logs after backup, select **Prune archive logs after backup** and select one of the following options:
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 a 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	Select Delete archive logs older than and 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

5. In the **Retention** page, specify the retention settings for the data backup and archive log backups:

If you want to...	Then...
Retain Snapshot copies based on retention settings specified in the storage system	<p>Select Keep Snapshot copies based on the storage systems maximum retention settings.</p> <p>If <code>autodelete</code> is set to NO in ONTAP, then after retaining 255 Snapshot copies, backup operation fails. If <code>autodelete</code> is set to YES, then after retaining the maximum limit of 255 Snapshot copies, ONTAP deletes the older Snapshot copies for new Snapshot copies.</p> <p>Note: Snapshot copies are deleted by ONTAP and therefore there might be backups in the SnapCenter that are not actually available for restore or cloning.</p>
Retain a certain number of Snapshot copies	<p>Select Total Snapshot copies to retain and specify the number of Snapshot copies that you want to retain.</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>Important: You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retentioncount 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>
Retain the Snapshot copies for a certain number of days	<p>Select Delete Snapshot copies older than and specify the number of days for which you want to retain the Snapshot copies before deleting them.</p>

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

Select a retention period for your backups

Data backup retention settings i

- Total Snapshot copies to keep
- Keep Snapshot copies for
- Keep Snapshot copies based on the storage systems maximum retention settings

Archive log backup retention settings

- Total Snapshot copies to keep
- Keep Snapshot copies for
- Keep Snapshot copies based on the storage systems maximum retention settings

6. 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).
Secondary policy label	Select a Snapshot label. Depending on the Snapshot label that you select, Data ONTAP 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.

7. 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.

8. In the **Schedule** page, select the frequency of the backup operation from the **Schedule** type drop-down list:

If you select...	Then...
OneTime	Specify the start date.

If you select...	Then...
Hourly	Specify the start date, the interval when the operation should occur, and optionally an expiry date.
Daily	Specify the start date, the days and interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.
Weekly	Specify the start date, the interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.
Monthly	Specify the start date, the month, days, and interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.

If you do not want to schedule the backup operation, you must select **None**.

- Review the summary and click **Finish**.

Related concepts

[How resources, datasets, and policies are used in data protection](#) on page 10

Related tasks

[Creating verification policies for Oracle databases](#) on page 22

Related information

[SnapCenter Software 1.1 Installation and Setup Guide](#)

Creating backup datasets and attaching policies for Oracle databases

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

Steps

- In the left navigation pane, click **Datasets**.
- Click **New > Backup Dataset**.
- In the **Name** page, perform the following:

For this field...	Do this...
Dataset name	Enter a name for the dataset.
Description	Enter information about the dataset.
Plug-in	Select SnapCenter Plug-in for Oracle Database .
Policy	Select the Oracle database specific policy that you want to associate with the dataset.

For this field...	Do this...
Windows scheduler Run as	Select windows run as credential that you had created for schedules. This option is available only if you selected a policy that has a schedule.
Use custom name format for Snapshot copy	Select this check box and enter the custom name format that you want to use for the Snapshot copy name. For example, <i>customtext_dataset_policy_hostname</i> or <i>dataset_hostname</i> . By default, a timestamp is appended to the Snapshot copy name.
Backup archive logs before missing archive log files	Select this check box to backup all archive log files except the missing archive log files. If you do not select this check box, only the archive log files that are created after the missing archive log files are backed up.
Exclude archive log destinations from backup	Specify the destinations of the archive log files that you do not want to back up.

- In the **Resources** page, select the Oracle database host name from the Host drop-down list.
Note: The database will be listed in the Available Resources section only if the database is discovered successfully.
- Select the databases from the **Available Databases** section and click the right arrow to move them to the **Selected Databases** section.
- In the **Verification** page, optionally select the SnapMirror or SnapVault location from where you want to verify the backup.
- In the **Notification** page, perform the following:

If you want to...	Do this...
Log SnapCenter Server events to the storage system's syslog	Select Log SnapCenter Server events to storage system syslog .
Send AutoSupport messages to technical support for failed operations	Select Send AutoSupport notification for failed operations to the storage system
Send emails	Select Email preference and specify the scenarios under which you want the emails to be sent. You must also specify the SMTP server name, the sender and receiver email addresses, and the subject of the mail.

- Review the summary and click **Finish**.

Backing up resources

You can either perform on-demand backup of a specific resource from the Inventory page or perform on-demand backup of a dataset containing multiple resources from the Datasets page.

About this task

If you want to schedule your backup, you can specify the schedule details while creating the backup policy. You can schedule multiple backups to run at the same time across multiple Oracle database hosts.

Note: While creating a backup, 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.

Choices

- [Backing up resources on demand](#) on page 29
If a resource is not yet part of any dataset, you can back up the resource from the Inventory page.
- [Backing up datasets on demand](#) on page 30
If you want to back up multiple resources simultaneously, you can back up the dataset containing multiple resources from the Datasets page.

Related tasks

[Creating backup policies for Oracle databases](#) on page 23

[Backing up Oracle databases using Linux commands](#) on page 67

Backing up resources on demand

If a resource is not yet part of any dataset, you can back up the resource from the Inventory page.

Before you begin

You must have created a backup policy.

About this task

A dataset is required to perform all backup jobs, even when you are backing up a resource. You will be prompted to create a dataset when initiating the backup.

Steps

1. In the left navigation pane, click **Inventory**.
2. In the **Inventory** page, select the database that you want to back up and click **Backup Now**.
3. In the **Backup Now** page, enter a name for the dataset, an optional description, and select a policy from the drop-down box.
4. Click **OK**.

Related tasks

[Creating backup policies for Oracle databases](#) on page 23

[Monitoring backup operations from the Jobs page](#) on page 30

Backing up datasets on demand

If you want to back up multiple resources simultaneously, you can back up the dataset containing multiple resources from the Datasets page.

Before you begin

You must have created a backup policy and a dataset.

Steps

1. In the left navigation pane, click **Datasets**.
2. In the **Datasets** page, select the dataset that you want to back up and click **Backup Now**.
3. In the **Backup** page, select a policy from the drop-down box, and click **Backup**.

Related tasks

[Creating backup policies for Oracle databases](#) on page 23

[Creating backup datasets and attaching policies for Oracle databases](#) on page 27

[Monitoring backup operations from the Jobs page](#) on page 30

Monitoring backup operations from the Jobs page

You can monitor the progress of different SnapCenter backup 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

Steps

1. In the left navigation pane, click **Monitor**.
2. In the top navigation bar, click **Jobs**.
3. In the **Jobs** page, filter the list so that only backup operations are listed:
 - a. Click **Filter**.
 - 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  to view the operations completed successfully.
4. Select the backup job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

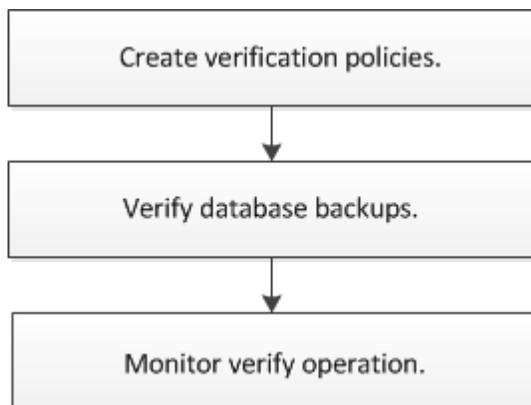
Verifying Oracle database backups

The verify workflow includes creating verification policies, performing the verify operation, and monitoring the operation. Verification runs database integrity checks on backups.

About this task

You do not have to configure a remote verification server to verify the Oracle database backups. When you verify a database backup, the host on which that database resides is used as the verification server. In a RAC environment, the preferred backup node is used as the verification server.

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



Creating verification policies for Oracle databases

You must create a verification policy to verify your backup. A verification policy is a set of rules that governs how you manage and schedule verification operations. Additionally, you can specify replication, prescript, and postscript details.

Steps

1. In the left navigation pane, click **Policies**.
2. Click **New > Oracle Policy > Verification**.
3. In the **Name** page, enter the policy name and description.
4. In the **Schedule** page, select the frequency of the verification operation from the **Schedule type** type drop-down list:

If you select...	Then...
OneTime	Specify the start date.
Hourly	Specify the start date, the interval when the operation should occur, and optionally an expiry date.
Daily	Specify the start date, the days and interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.
Weekly	Specify the start date, the interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.

If you select...	Then...
Monthly	Specify the start date, the month, days, and interval when the operation should occur, and optionally a recurrence in minutes and/or an expiry date.

If you do not want to schedule the verification operation, you must select **None**.

5. In the **Replication** page, select **Verify on secondary storage** if you want to verify the backup on a secondary storage volume.

The Snapshot copies are verified on the secondary storage volume and the primary Snapshot copies are marked as verified.

You must ensure that the SVM settings of the secondary storage are included in the SnapCenter Server settings if you are verifying on the secondary storage.

If you want to verify the backups soon after the backup operation is completed, you must select the secondary replication option while creating the backup policy.

6. 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 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.

7. In the **Options** page, enter the number of most recent backups you want to verify.

Example

If three backups were created as part of the schedule and you enter 2 as the number of backups you want to verify; then of the three backups, 2 most recent backups will be verified.

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

Related concepts

[How resources, datasets, and policies are used in data protection](#) on page 10

Related information

[SnapCenter Software 1.1 Installation and Setup Guide](#)

Verifying an Oracle database backup

The backup verification process protects you from restoring a backup that contains any physical-level corruption or missing database tables or rows. Verifying backups ensures that the databases can be restored as required.

Before you begin

- You must have created a verification policy.
You can either verify backups on demand or schedule the verification jobs. If you want to schedule the verification jobs, you must create the verification policy with a schedule.
- You must have backed up a dataset to verify.

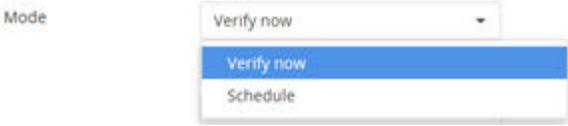
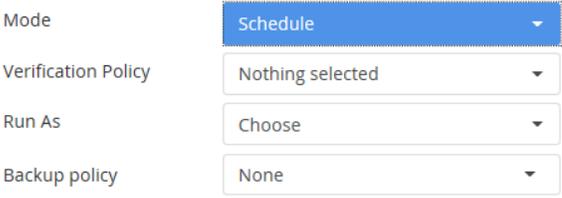
- You must have an unverified backup to verify.
- If you have an Automatic Storage Management (ASM) database instance in NFS environment and want to verify the ASM backups, you must have added the ASM disk path `/var/opt/snapcenter/sco/backup_*/**/**/*/*` to the existing path defined in the `asm_diskstring` parameter.

About this task

You can verify only the backed up datafiles.

Steps

1. In the left navigation pane, click **Datasets**.
2. Select the backup dataset and click **Verify**.
3. In the **Verification** page, perform the following:

For this field...	Do this...
Mode	<ul style="list-style-type: none"> • Select Verify now if you want to verify backups on demand. • Select Schedule if you want to verify backups as per the schedule defined in verification policy. 
Verification Policy	Select the verification policy.
Run As	<p>Select the Windows run as credentials that must be used for scheduled verifications.</p> <p>This option is displayed only if you have selected the verify mode as Schedule.</p> 
Backup policy (Optional)	Select the backup policy associated with the dataset.

4. Click **OK**.

Monitoring verification operations from the Jobs page

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

About this task

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

-  In progress

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

Steps

1. In the left navigation pane, click **Monitor**.
2. In the top navigation bar, click **Jobs**.
3. In the **Jobs** page, filter the list so that only verification operations are listed:
 - a. Click **Filter**.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Verification**.
 - d. From the **Status** drop-down list, select the backup status.
 - e. Click  .
4. Select the verification job, and then click **Details** to view the job details.
You can view the verification status by clicking **Manage Backups**.
5. In the **Job Details** page, click **View logs**.

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.

Related tasks

[Mounting a database backup](#) on page 36

[Unmounting a database backup](#) on page 37

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.

Steps

1. In the left navigation pane, click **Inventory**.
2. Select **Oracle Database**.
3. To filter the list of resources, select the host from the Host drop-down list, and then select the database type from the Database Type drop-down list.
4. Select the resource and click **Mount**.
5. In the **Mount backups** page, perform the following:
 - a. Select whether you want to mount from the primary or secondary storage systems, and then select the backup to mount.

To filter the list of backups, you can either specify the backup name in the **search** field or use the **Filter** drop-down list to specify the time range.
 - b. From the **Choose the host to mount the backup** drop-down list, select the host on which you want to mount the backup.

The mount path `/var/opt/snapcenter/sco/backup_mount/backup_name/database_name` is displayed.
6. Click **Mount**.

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

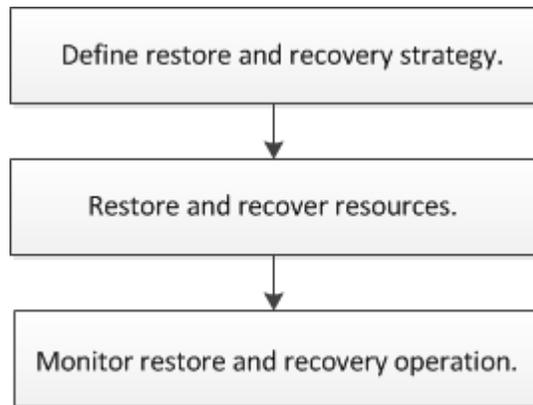
1. In the left navigation pane, click **Inventory**.
2. Select **Oracle Database**.
3. To filter the list of resources, select the host from the Host drop-down list, and then select the database type from the Database Type drop-down list.
4. Select the resource and click **Unmount**.
5. In the **Unmount backups** page, select a backup to unmount.
To filter the list of mounted backups, you can either specify the backup name in the **search** field or use the **Filter** drop-down list to specify the time range.
6. Click **Unmount**.

Restoring and recovering Oracle databases

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

About this task

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



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

[SnapCenter Software 1.1 Linux Command Reference Guide](#)

Related concepts

[Backing up, restoring, and cloning using Linux commands](#) on page 67

Defining a restore and recovery strategy for Oracle databases

You must define a strategy before you restore and recover your database so that you can perform restore and recover operations successfully.

Steps

1. Determine the backups that can be used for restore and recovery operations.
2. Know the limitations that are related to the restore and recovery operations.
3. Determine the restore methods that are supported.
4. Decide the type of restore and recovery operations that you want to perform.
5. Determine the source and destination for the restore operation.

Related references

[Backups supported for restore and recovery operations](#) on page 39

Limitations related to restore and recovery operations on page 40

Types of restore methods supported for Oracle databases on page 40

Types of restore operations supported for Oracle databases on page 40

Types of recovery operations supported for Oracle databases on page 41

Sources and destinations for a restore operation on page 41

Backups supported for restore and recovery operations

Depending on your version of Oracle Database, SnapCenter supports restore and recovery of different types of backups.

- Oracle 11g
 - Standalone online data backup
 - Standalone offline shutdown data backup
 - Standalone offline mount data backup
 - Standalone full backup
 - Offline mount backup of databases in a Data Guard configuration
 - Backup of databases in an Active Data Guard configuration
 - Online data, online full, offline mount, and offline shutdown backups in Real Application Clusters (RAC) configuration
 - Online data, online full, offline mount, and offline shutdown backups in Automatic Storage Management (ASM) configuration
- Oracle 12c (legacy)
 - Standalone online data backup
 - Standalone offline shutdown data backup
 - Standalone offline mount data backup
 - Standalone full backup
 - Offline mount backup of databases in a Data Guard configuration
 - Backup of databases in an Active Data Guard configuration
 - Online data, online full, offline mount, and offline shutdown backups in RAC configuration
 - Online data, online full, offline mount, and offline shutdown backups in ASM configuration
- Oracle 12c (CDB)
 - Online data backup
 - Offline shutdown data backup
 - Offline mount data backup
 - Full backup
 - Offline mount backup of databases in a Data Guard configuration
 - Backup of databases in an Active Data Guard configuration
 - Online data, online full, offline mount, and offline shutdown backups in RAC configuration

- Online data, online full, offline mount, and offline shutdown backups in ASM configuration

Limitations related to restore and recovery operations

Before you perform restore and recovery operations, you must be aware of the limitations.

The following restore and recovery operations are not supported:

- Restore and recovery of tablespaces of the root container database (CDB)
- Restore of temporary tablespaces and temporary tablespaces associated with PDBs
- Restore and recovery of tablespaces from multiple PDBs simultaneously
- Restore of log backups
- Restore of backups to a different location
- Partial restore with recovery up to a specific SCN and date
- Restore from secondary storage if one of the volumes do not have a SnapMirror or SnapVault relationship for a given database.
- Restore of redo log files in any configuration other than Data Guard or Active Data Guard.

Types of restore methods supported for Oracle databases

SnapCenter supports connect-and-copy or in-place restore for Oracle databases. During a restore operation, SnapCenter determines the restore method appropriate for the file system to be used for restore without any data loss.

In-place restore

If the database layout is similar to the backup and has not undergone any configuration change on the storage and database stack, in-place restore is performed, wherein the restore of file or LUN is performed on Data ONTAP.

Note: Data ONTAP 8.3 or later supports in-place restore from secondary location.

Connect-and-copy restore

If the database layout differs from the backup or if there are any new files after the backup was created, connect-and-copy restore is performed. In the connect-and-copy restore method, the following tasks are performed:

1. The LUN or volume is cloned from the Snapshot copy and the file system stack is built on the host using the cloned LUNs or volumes.
2. The files are copied from the cloned file systems to the original file systems.
3. The cloned file systems are then unmounted from the host and the cloned volumes are deleted from Data ONTAP.

Types of restore operations supported for Oracle databases

SnapCenter enables you to perform different types of restore operations for Oracle databases.

Before restoring the database, backups are validated to identify whether any files are missing when compared to the actual database files.

Full restore

- Restores only the datafiles

- Restores only the control files
- Restores the datafiles and control files
- Restores datafiles, control files, and redo log files in Data Guard and Active Data Guard configurations

Partial restore

- Restores only the selected tablespaces
- Restores only the selected pluggable databases (PDBs)
- Restores only the selected tablespaces of a PDB

Types of recovery operations supported for Oracle databases

SnapCenter enables you to perform different types of recovery operations for Oracle databases.

- The database up to the last transaction (all logs)
- The database up to a specific system change number (SCN)
- The database up to a specific date and time
You must specify the date and time for recovery based on the database host's time zone.

SnapCenter also provides the No recovery option for Oracle databases.

Note: The plug-in for Oracle database does not support recovery if you have restored using a backup that was created with the database role as standby. You must always perform manual recovery for physical standby databases.

Sources and destinations for a restore operation

You can restore an Oracle database from a backup copy on either primary or secondary storage. You can only restore databases to the same location on the same database instance.

Sources for a restore operation

You can restore databases from a backup on primary or secondary storage. If you want to restore from a backup on the secondary storage in a multiple mirror configuration, you can select the secondary storage mirror as the source.

Destinations for a restore operation

You can only restore databases to the same location on the same database instance.

In a Real Application Cluster (RAC) setup, you can restore RAC databases from any nodes in the cluster.

Restoring and recovering 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 recover the database. Recovery operation is performed using archive logs of the database in an active file system.

Before you begin

- You must have defined your restore and recovery strategy.

- The SnapCenter administrator must have assigned you the Storage Virtual Machines (SVMs) for both the source and destination volumes if you are replicating Snapshot copies to a mirror or vault.
- If archive logs are pruned as part of backup, you must have manually mounted the required archive log backups.
- If you want to restore Oracle databases residing on VMDK devices, you must ensure that the guest machine has the required number of free slots for allocating cloned VMDK disks.

About this task

When you restore a database, 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 restored, the operational lock file is automatically removed.

Steps

1. In the left navigation pane, click **Inventory**.
2. Select **Oracle Database**.
3. To filter the list of resources, select a host from the Host drop-down list, and then select the database type from the Database Type drop-down list.
4. Select the resource and click **Restore**.
5. In the **Backups** page, select whether you want to restore resources from the primary or the secondary storage system, and then select the backup copy to restore.

If you are restoring a RAC database, from the Select RAC node drop-down list, select the RAC node.

6. In the **Restore Scope** page, 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 or Active Data Guard configurations.
Pluggable databases (PDB)	Select Pluggable databases and specify the PDBs that you want to restore.
Pluggable database (PDB) tablespaces	Select Pluggable database (PDB) tablespaces and 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.

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

If you...	Do this...
Want to recover to the last transaction	Select All Logs .

If you...	Do this...
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 specify the location of external archive log files. If archive logs are pruned as part of backup and you have manually mounted the required archive log backups, specify the mounted backup path as the external archive log location for recovery.

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

Choose Recovery Scope

All Logs

Until SCN (System Change Number)

Date and Time




Date-time format: MM/DD/YYYY hh:mm:ss

No recovery

Specify external archive log files locations  

8. In the **PreOps** page, perform the following:
 - a. 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 `/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.
 - b. 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 it must be changed to a lower state to perform a restore operation. If the database is in a lower state, but 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 database is in open state and for restore it needs to be in mounted state, then the database state will be changed only if you have selected this check box.

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

- 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 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.

Note: After restoring either user tablespace (PDB) with control files, 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 those PDBs.

10. In the **Notification** page, perform the following:

If you want to...	Do this...
Log SnapCenter Server events to the storage system's syslog	Select Log SnapCenter Server events to storage system syslog .
Send AutoSupport messages to technical support for failed operations	Select Send AutoSupport notification for failed operations to the storage system
Send emails	Select Email preference and specify the scenarios under which you want the emails to be sent. You must also specify the SMTP server name, the sender and receiver email addresses, and the subject of the mail.

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

Related tasks

[Mounting a database backup](#) on page 36

[Restoring and recovering Oracle databases using Linux commands](#) on page 68

Monitoring restore operations from the Jobs page

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 database recovery states describe the conditions of the database 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

Steps

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

Cloning Oracle database backups

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

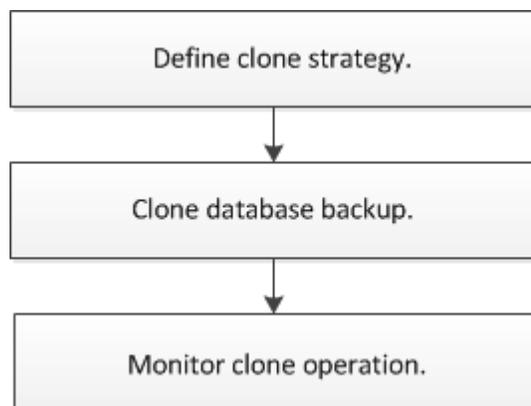
About this task

You might clone database backups for the following reasons:

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

You can clone a database backup residing on primary or secondary storage to the same host as that of the source database or to an alternate host with the same Oracle and OS version.

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



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

[SnapCenter Software 1.1 Linux Command Reference Guide](#)

Related concepts

[Backing up, restoring, and cloning using Linux commands](#) on page 67

Defining a clone strategy for Oracle databases

Defining a strategy before cloning your database ensures that the cloning operation is successful.

Steps

1. Determine the backups that can be used for cloning.
2. Decide the type of cloning that you require.

Related references

[Backups supported for cloning](#) on page 47

[Type of cloning supported](#) on page 47

Backups supported for cloning

SnapCenter supports cloning of different backups of Oracle databases.

- Oracle 11g, Oracle 12c (legacy), and Oracle 12c (CDB)
 - Online data backup
 - Online full backup
 - Offline mount backup
 - Offline shutdown backup
 - Backup of databases in a Data Guard and Active Data Guard configurations
 - Online data, online full, offline mount, and offline shutdown backups in Real Application Clusters (RAC) configuration
 - Online data, online full, offline mount, and offline shutdown backups in Automatic Storage Management (ASM) configuration

Note: Archive log backups are not supported for cloning.

Type of cloning supported

In an Oracle Database environment, SnapCenter supports cloning of a database backup. You can clone the backup from primary and secondary storage systems.

The SnapCenter Server uses NetApp FlexClone technology to clone backups.

Cloning an Oracle database backup

You can use SnapCenter to clone a database backup.

Before you begin

- You must have created a database backup.
- If you want to clone a backup of 12c database, you must have set the value of `exclude_seed_cdb_view` to **FALSE** to retrieve seed PDB related information.
The seed PDB is a system-supplied template that the CDB can use to create new PDBs. The seed PDB is named PDB\$SEED and is required for cloning a backup of 12c database. For information about PDB\$SEED, see the Oracle Doc ID 1940806.1.

About this task

While cloning a backup of an ASM database in 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.

While cloning from a RAC environment to standalone, the `_no_recovery_through_resetlogs` parameter is added to the `initcloneSID.ora` file. After successful cloning and recovery, you can manually remove the `_no_recovery_through_resetlogs` parameter.

Steps

1. In the left navigation pane, click **Inventory**.
2. Select **Oracle Database**.
3. To filter the list of resources, select the host from the Host drop-down list, and then select the database type from the Database Type drop-down list.
4. Select the resource and click **Clone Backup**.
5. In the **Source** page, perform the following:
 - a. In the **Clone SID** field, enter the SID of the clone.
The maximum length of the clone SID is 8 characters.
 - b. Select whether you want to clone the backup from the primary or secondary storage system.
 - c. Select the backup that you want use for cloning and click **Next**.
6. In the **Locations** page, perform the following:

For this field...	Do this...
Clone host	By default, the host is populated. If you want to specify a different host, select the host on which you want to clone the backup.
Datafile locations	By default, the datafile location is populated. If you want to specify a different path, enter the datafile mount points or ASM disk group names for clone database.
Control files	By default, the control file path is populated. If you want to specify a different path, enter the control file path.
Redo logs	By default, the redo log file group, path, and their sizes are populated. If you want to change the default values, enter the redo log file group, path, and their sizes.

7. In the **Credentials** page, perform the following:

For this field...	Do this...
Run As for sys user	Select the Run As account to be used for defining the sys user password of the clone database. If you do not want to specify the run as account, you must select None . The OS authentication will be used if you do not specify the run as account. The default value is None .
ASM Instance Run As	Select the Run As account of the ASM instance on the alternate host. This option is available only if you have configured ASM. If you do not want to specify the run as account, you must select None . The OS authentication will be used if you do not specify the run as account. The default value is None .

The Oracle home, user name, and group details are automatically populated from the backup metadata. However, you can change the automatically populated values.

Database Credentials for the clone

Run As for sys user	<input type="text" value="None"/>	i
ASM Instance Run as	<input type="text" value="None"/>	i
ASM Port	<input type="text" value="1521"/>	

Oracle Home Settings

Oracle Home	<input type="text"/>
Oracle OS User	<input type="text" value="Enter the oracle user"/>
Oracle OS Group	<input type="text" value="Enter the oracle group"/>

8. Optional: In the **PreOps page, perform the following:**

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

You must store the prescripts 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. In the **Database Parameter settings** section, you can modify the values of the default database parameters and add custom parameters.

These parameters are used to initialize the database.

Note: The default value of `log_archive_dest_1` is `$ORACLE_HOME/clone_sid` and the archive logs are created in this location. If you have deleted the `log_archive_dest_1` parameter, the archive log location is determined by Oracle.

[Database Parameter settings](#)

audit_file_dest	/ora01/app/oracle/admin/tst/adump	<input type="text" value="x"/>	<input type="text" value="↑"/>	<input text"="" type="text" value="x"/>		
log_archive_dest_1	LOCATION=/ora01/app/oracle/product/12c/db...	<input type="text" value="x"/>				
log_archive_format	%t_%s_%r.dbf	<input type="text" value="x"/>	<input type="text" value="↓"/>	<input type="text" value="Reset"/>		

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

For this field...	Do this...
Recover Database	<p>Select this option if you want to perform recovery after cloning. If you selected recovery, you must also define the recovery scope by specifying one of the following:</p> <ul style="list-style-type: none"> • Select Date and Time to recover the database up to a specific data and time. • Select Until SCN (System Change Number) to recover the database up to a specific system change number (SCN). • Select No recovery if you do not want to perform recovery. This option is not available if you are trying to clone a backup of a Data Guard or an Active Data Guard database. <p>Note: Recovery is not supported for Data Guard and Active Data Guard databases. In a Data Guard or Active Data Guard configurations, you must perform a manual recovery after the clone operation is completed and then open the database with <code>resetlogs</code>.</p> <p>Recovery point is calculated based on the latest log backups available before pruning. If the recovery point has to be more recent, you must manually mount the required log backups and specify the log path in the Specify external archive log locations field.</p> <p>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.</p>
Enter SQL statements to apply when clone is created	Enter the SQL statements that you want to apply after the clone is created.
Enter scripts to run after clone operation	<p>Enter the path and the arguments of the postscript that you want to run after the clone operation.</p> <p>You must store the postscripts either in <code>/var/opt/snapcenter/spl/scripts</code> or in any folder inside this path. By default, the <code>/var/opt/snapcenter/spl/scripts</code> path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.</p>

10. In the **Notification** page, perform the following:

If you want to...	Do this...
Log SnapCenter Server events to the storage system's syslog	Select Log SnapCenter Server events to storage system syslog .
Send AutoSupport messages to technical support for failed operations	Select Send AutoSupport notification for failed operations to the storage system
Send emails	<p>Select Email preference and specify the scenarios under which you want the emails to be sent.</p> <p>You must also specify the SMTP server name, the sender and receiver email addresses, and the subject of the mail.</p>

11. Review the summary and 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 use this clone to perform manual recovery.

Related tasks

[Cloning Oracle database backups using Linux commands](#) on page 69

Monitoring clone operations from the Jobs page

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

Steps

1. In the left navigation pane, click **Monitor**.
2. In the top navigation bar, click **Jobs**.
3. In the Jobs page, filter the list so that only clone operations are listed:
 - a. Click **Filter**.
 - 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 .
4. Select the clone job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

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 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**.

Backup operation fails when initiated before starting the database instance

Description

After starting the Automatic Storage Management (ASM) database instance, if you initiate an offline shutdown backup before starting the database instance, the backup operation fails.

Error message

```
ORACLE-00507: Failed to find entry for SID SID or database name
database_name in /etc/oratab file on host_name.
```

Corrective action

The backup operation fails because the database instance details are missing from the `/etc/oratab` file. You must manually enter the database instance details in the `/etc/oratab` file and then start the database instance.

Backup fails with license issue

Description

Backup operation fails because the SnapManager suite license is not installed in the clustered Data ONTAP storage system.

Error message

```
'SnapCenter Plug-in for Oracle License' failed with error: No storage
license found. Please contact administrator.
```

Corrective action

Install the SnapManager suite license on the controller.

Backup fails during the file system discovery on a VM

Description

When you perform a VMDK backup on a virtual machine (VM), backup operation fails while discovering the file system.

Error message

Failed to retrieve the unit serial number for the device '/dev/sde'.

Corrective action

You must set the `disk.enableUUID true` configuration parameter on the VM under edit settings **VM options > Advanced > Edit configuration**.

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

Unable to find any Storage Virtual Machine (SVM) to verify the plugin license. Please check your SVM is registered with SnapCenter.

Corrective action

You must ensure that you are using the supported storage environment. SnapCenter supports only clustered Data ONTAP storage systems.

Registering backup activity fails during the backup operation

Description

Registering backup activity fails during the backup operation due to the mismatch of the host name between SnapCenter and the Linux host.

Corrective action

1. Check the DNS entry in `/etc/hosts` on the Linux host (if any).
Example: `10.230.156.44 scspr0090825001.gdl.englab.netapp.com`
`scspr0090825001`
2. Verify the DNS entry on the Linux host: `nslookup`
3. Verify the entry in `C:\Windows\System\drivers\etc\hosts` on the SnapCenter Server (if any).
4. Verify the DNS entry on the SnapCenter Server: `nslookup`
5. If there is a mismatch between the two DNS entries, then change the host name format from short name to FQDN.
Ensure that the `hostname` command in Linux returns the short name and the `hostname -f` returns FQDN.
6. Restart the SnapCenter Plug-in Loader (SPL) for the changes to take effect.

Verification of datafiles backup fails

Description

If you have disabled OS authentication for an Automatic Storage Management (ASM) database, and when you try to verify the datafiles backup of that ASM database, the operation fails.

Corrective action

You must enable OS authentication for the ASM database.

Backup verification fails

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

You must perform the following:

1. Increase the value of the `VERIFICATION_DELAY` and `VERIFICATION_RETRY_COUNT` parameter in `sco.properties` file located at `/var/opt/snapcenter/sco/etc`.
2. Restart the SnapCenter Plug-in Loader (SPL) service.

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.

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

[SnapCenter Software 1.1 Installation and Setup Guide](#)

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 volume fails

Description

Restore operation from a secondary SnapMirror or SnapVault volume fails if you have configured load-sharing mirror (LSM) 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 large number of backup copies on the primary volume so that the restore operation can be performed from the primary volume.
- Mount the backup from the secondary storage and manually copy the files that are to be restored.

Cloning operation on the primary storage fails

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 configuration stack in the Linux host.

Related information

[Linux Unified Host Utilities 7.0 Quick Start Guide](#)

Cloning operation fails in SAN environments in OL 7 or later or RHEL 7 or later

Description

If you are using Oracle Linux 7 or later or Red Hat Enterprise Linux (RHEL) 7 or later, cloning operation fails in storage area network (SAN) environments. This issue occurs because by default, lvm2-lvmetad service is enabled.

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.

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:

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 new backup.
3. Perform cloning using the new backup.

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 in `sco.properties` file located at `/var/opt/snapcenter/sco/etc`.

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.

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

Clustered Data 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_dbssid`) 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.

Error message

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

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.
```

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 `PERL_COMMAND_EXECUTION_TIMEOUT` parameter in the `scu.properties` file located at `/var/opt/snapcenter/scu/etc/`.

The `PERL_COMMAND_EXECUTION_TIMEOUT` parameter specifies the number of seconds to wait for an operation to complete. The default value is 1800 seconds.

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 Oracle database datasets

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

About this task

You can perform the following tasks related to datasets:

- Create a backup dataset.
- Modify a backup dataset.
- Create a backup using the dataset.
- Verify a backup using the dataset.
- Delete a backup dataset.

Related tasks

[Creating backup datasets and attaching policies for Oracle databases](#) on page 27

[Backing up datasets on demand](#) on page 30

If you want to back up multiple resources simultaneously, you can back up the dataset containing multiple resources from the Datasets page.

[Verifying an Oracle database backup](#) on page 33

Modifying datasets

You can modify a dataset to edit the information that you provided when creating the dataset.

Steps

1. In the left navigation pane, click **Datasets**.
2. Select a dataset and click **Modify**.
3. Edit the information and click **Finish**.

Stopping operations on datasets temporarily

You can temporarily stop the operations that are being performed on a dataset.

Steps

1. In the left navigation pane, click **Datasets**.
2. Select the dataset, click **Maintenance**, and then click **OK**.

Resuming operations on datasets

You can resume the operations on a dataset that you had stopped temporarily.

Steps

1. In the left navigation pane, click **Datasets**.
2. Select the dataset, click **Production**, and then click **OK**.

Deleting datasets for Oracle databases

You can delete a dataset if you no longer need it to perform different operations. You must ensure that datasets are deleted before you remove plug-ins from SnapCenter.

About this task

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

Steps

1. In the left navigation pane, click **Datasets**.
2. Select the dataset and click **Delete**.
3. Select the **Delete backups and policies associated with this dataset** check box to remove all backups, metadata, policies, and Snapshot copies associated with the dataset.
4. Click **OK**.

Managing Oracle database policies

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

About this task

You can perform the following tasks related to policies:

- Create a backup or verification policy.
- Modify a backup or verification policy.
- Copy a backup or verification policy.
- View details of a backup or verification policy.
- Delete a backup or verification policy.

Related tasks

[Creating backup policies for Oracle databases](#) on page 23

[Creating verification policies for Oracle databases](#) on page 22

Modifying policies

You can edit a policy to modify the information that you provided when creating the policy. You might want to change the replication options, Snapshot copy retention settings, scripts, and types.

Steps

1. In the left navigation pane, click **Policies**.
2. Select the policy and click **Modify**.
3. Modify the information and click **Finish**.

Copying policies

You can copy policies if you want to create a policy. Copying a policy rather than creating a new one saves time.

Steps

1. In the left navigation pane, click **Policies**.
2. Select the policy and click **Copy**.
3. Accept the default name or type a new name and click **OK**.

Viewing policy details

You can view details of a policy before you perform any operation using that policy or create a copy of that policy. You might want to view the details to ensure that the correct options will be applied to the operation.

Steps

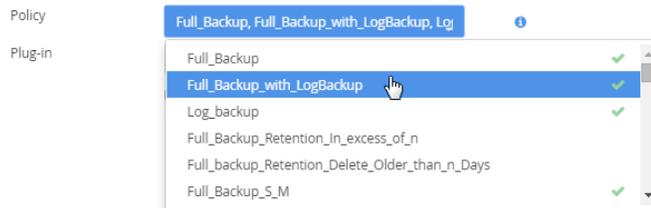
1. In the left navigation pane, click **Policies**.
2. Select the policy and click **Details**.
You can also double-click the policy to view the details of the policy.
3. Review the details and click **Close**.

Detaching policies from a dataset

Any time that you no longer want policies for a dataset, you can detach them.

Steps

1. In the left navigation pane, click **Datasets**.
2. Select the dataset and click **Modify**.
3. In the **Name** page of the **Modify Dataset** wizard, clear the check mark next to the names of the policies you want to detach:



4. Make any additional modifications to the dataset in the rest of the wizard and 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 datasets.

Steps

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

Managing backups

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

Renaming or deleting backup copies

You can rename or delete backup copies of a selected database. If the backup is associated with a cloned database, you cannot delete the backup copy.

Before you begin

- You must have deleted the associated clones.

Steps

1. In the left navigation pane, click **Inventory**.
2. Select the SnapCenter plug-in that you are using to manage your resources.
3. To filter the list, select the host from the **Host** drop-down list, and then select the database type from **Database Type** drop-down list.
4. Select the resource and click **Manage Backups**.
5. On the **Manage Backups** page, select the backup copy and choose one of the following options:

Option	Do this...
Renaming the backup	In the Rename as field, enter a new name and click Rename .
Deleting the backup	Click Delete > OK .

If you want to delete multiple backups, select the backups and click **DeleteOK**.

6. Click **Close**.

Related information

[SnapCenter Software 1.1 Linux Command Reference Guide](#)

Managing clones

Using SnapCenter, you can view details about the clones you have created and delete them if you find them no longer necessary.

Viewing clone details

You can view details about the clones associated with a database. You might want to view clone details for review or to ensure that the correct clone was deleted.

Steps

1. In the left navigation pane, click **Inventory**.
2. Select the SnapCenter plug-in that you are using to manage your resources.
3. To filter the list of resources, select the host from the Host drop-down list.
4. Select the resource from which the clone was created and click **Manage Clone** to view the list of clones with details taken from the parent database.
5. Click **Close**.

Deleting clones

You can delete clones associated with a database 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* clone and then delete *db1_clone1*.

Steps

1. In the left navigation pane, click **Inventory**.
2. Select the SnapCenter plug-in that you are using to manage your resources.
3. To filter the list, select the host from the **Host** drop-down list.
4. Select the resource from which the clone was created and click **Manage Clone** to view the list of clones.
5. Select the clone and click **Delete**.
6. Click **OK** to confirm the deletion.

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 1.1 Linux Command Reference Guide](#)

Backing up Oracle databases using Linux commands

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

Before you begin

- You must have added the Storage Virtual Machine (SVM) connection 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 backup dataset, backup and verify the backup.

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

[SnapCenter Software 1.1 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 or verification 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 dataset.
6. Add a dataset to the SnapCenter: `Add-SmBackupDataset`
7. Initiate a new Snapshot copy job: `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`
9. Initiate the verification job by specifying the dataset backup that you want to verify and the verification policy for the operation: `Invoke-SmBackupVerification`

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 1.1 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 1.1 Linux Command Reference Guide](#)

Steps

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

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 dataset 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`.

Where to go next

You can find more information about different features and release-specific information for SnapCenter in the documentation available on the NetApp Support Site at mysupport.netapp.com.

- [*SnapCenter Software 1.1 Release Notes*](#)
Provides important information about this release of Snap Creator Server and the SnapCenter plug-in packages, including fixed issues, known issues, cautions, limitations, and any documentation updates or corrections.
- [*SnapCenter Software 1.1 Installation and Setup Guide*](#)
Describes the steps required to prepare for installation and to install SnapCenter and the SnapCenter plug-in packages. Setup processes are described for both current SnapManager users who are migrating to SnapCenter and users who are implementing a new SnapCenter environment.
- [*SnapCenter Software 1.1 Getting Started Guide*](#)
Provides simple information about how to get started with SnapCenter, regardless of which plug-in packages you have installed, and outlines a simple workflow to help you perform your first backup operation.
- [*SnapCenter Software 1.1 Administration Guide*](#)
Provides information about how to administer SnapCenter, provision Windows hosts with storage, configure and maintain role-based access control (RBAC), and use the centralized reporting options.
- [*SnapCenter Software 1.1 Windows Cmdlet Reference Guide*](#)
Provides reference information about the Windows PowerShell cmdlets available in SnapCenter, including a description of each cmdlet, its syntax, and examples for its use. This content is also available through the SnapCenter PowerShell cmdlet help.
- [*SnapCenter Software 1.1 Linux Command Reference Guide*](#)
Provides reference information about the Linux commands available for Linux plug-ins, including a description of each command, its syntax, and examples for its use. This content is also available through the SnapCenter command-line interface help.

Copyright information

Copyright © 1994–2016 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<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*](mailto: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., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- administrator role
 - logging in as [15](#)
- architecture
 - overview of SnapCenter components [7](#)
- archiving
 - using SnapVault [20](#)
- automation
 - scripts for [21](#)

B

- backing up
 - datasets [30](#)
 - individual resources [29](#)
 - multiple resources [30](#)
- backing up resources
 - datasets [30](#)
 - naming convention for copies [19](#)
 - on demand [29](#)
 - preparing for [13](#)
 - scheduled [29](#)
 - scheduling verification [20](#)
 - strategy [17](#)
 - workflow for [12](#)
- backup
 - defining a strategy [17](#)
 - identifying database resources [21](#)
 - Oracle database configurations supported for backup [17](#)
- backup and cloning fails
 - disk group is dismounted [57](#)
- backup copies
 - deleting [65](#)
 - managing [65](#)
 - number of [20](#)
 - renaming [65](#)
 - verifying [21](#)
- backup datasets
 - creating for Oracle databases [27](#)
 - managing [61](#)
- backup fails
 - file system discovery [52](#)
- backup failure
 - database instance not started [52](#)
 - license issue [52](#)
- backup operation
 - fails [53](#)
- backup operations
 - resources, datasets, and policies overview [10](#)
 - using Linux commands for Oracle databases [67](#)
 - workflow for Oracle databases [16](#)
- backup operations, SnapCenter
 - monitoring progress from the Jobs page [30](#)
- backup policies
 - creating for Oracle databases [23](#)
 - managing [63](#)
- backup types
 - modes [18](#)
 - offline backup described [18](#)
 - online backup described [18](#)

- backup verification
 - fails [54](#)
- backup, resources
 - number of backup copies [20](#)
 - number of days [20](#)
 - when to perform [19](#)
- backups
 - cloning databases from [47](#)
 - creating verification policies for Oracle databases [22](#), [32](#)
 - managing [65](#)
 - supported for cloning [47](#)
 - supported for restore and recovery operations [39](#)
 - types of [18](#)
 - verifying for Oracle database [33](#)
- backups, database
 - mounting [36](#)
 - supported cloning [47](#)
 - unmounting [36](#), [37](#)
- backups, Oracle database
 - workflow for cloning [46](#)

C

- clone delete
 - file system not deleted [57](#)
- clone details
 - view [66](#)
- clone operations
 - resources, datasets, and policies overview [10](#)
 - using Linux commands for Oracle databases [69](#)
 - workflow for Oracle database backups [46](#)
- clone operations, SnapCenter
 - monitoring progress from the Jobs page [51](#)
- clones
 - defining a strategy [46](#)
 - deleting [66](#)
 - managing [66](#)
 - supported type [47](#)
- cloning databases
 - from a backup [47](#)
 - preparing for [13](#)
- cloning fails
 - in SAN environments [56](#)
 - on primary storage [56](#)
- cloning Oracle databases
 - strategy [46](#)
- command execution timeout
 - error [59](#)
- commands, Linux
 - common tasks for scripting backup, restore, recovery, and clone operations [67](#)
- comments
 - how to send feedback about documentation [73](#)
- copying
 - policies [63](#)

- creating
 - backup datasets for Oracle databases [27](#)
 - backup policies for Oracle databases [23](#)
 - backups [29](#)
 - verification policies for Oracle databases [22, 32](#)

D

- data protection
 - environment setup [13](#)
 - resources, datasets, and policies overview [10](#)
- data protection operation
 - fails [59](#)
- database backups
 - mounting [36](#)
 - supported cloning [47](#)
 - unmounting [36, 37](#)
- database state
 - fails to change [55](#)
 - required for backup in RAC setup [19](#)
- databases
 - backing up on demand [29](#)
 - cloning from a backup [47](#)
 - destinations for a restore operation [41](#)
 - restoring and recovering [41](#)
 - sources for a restore operation [41](#)
- databases, Oracle
 - creating backup datasets for [27](#)
 - creating backup policies for [23](#)
 - creating verification policies for [22, 32](#)
 - defining a clone strategy [46](#)
 - identifying resources for backup [21](#)
 - policies [27](#)
 - supported configurations for backup [17](#)
 - verify workflow [32](#)
 - when to back up [19](#)
 - workflow for backing up [16](#)
 - workflow for cloning backups [46](#)
 - workflow for restoring and recovering [38](#)
- datafiles and control files
 - restore fails [55](#)
- datafiles backup
 - verification fails [54](#)
- datasets
 - backing up on demand [30](#)
 - deleting [62](#)
 - detaching policies from [64](#)
 - modifying [61](#)
 - operations
 - resuming on datasets [62](#)
 - resuming operations on [62](#)
 - temporarily stopping operations on [61](#)
 - used in SnapCenter data protection [10](#)
- datasets, backup
 - attaching policies [27](#)
 - creating for Oracle databases [27](#)
 - managing [61](#)
- deleting
 - backup copies [65](#)
 - clones [66](#)
 - datasets [62](#)
 - policies [64](#)
- disaster recovery

- using SnapMirror [20](#)
- disk paths
 - not included [54](#)
- display
 - incorrect time zone in log messages [59](#)
- documentation
 - additional resources [70](#)
 - how to receive automatic notification of changes to [73](#)
 - how to send feedback about [73](#)

F

- failure
 - database state change [55](#)
- failure, backup
 - database instance not started [52](#)
- failure, backup and clone operation
 - disk group is dismantled [57](#)
- failure, cloning operation
 - on primary storage [56](#)
 - SAN environments in OL 7 or later [56](#)
 - SAN environments in RHEL 7 or later [56](#)
- failure, recovery operation
 - of a cloned database [56](#)
- features
 - Plug-in for Oracle [6](#)
- feedback
 - how to send comments about documentation [73](#)
- file inaccessible
 - verification [54](#)
- file names
 - choosing for backup copies [19](#)
- file systems
 - not deleted [57](#)
- flowcharts
 - data protection workflow [12](#)
 - Oracle database restore and recovery workflow [38](#)
 - Oracle database verify workflow [32](#)
 - workflow for backing up Oracle databases [16](#)
 - workflow for cloning Oracle database backups [46](#)

H

- host name
 - FQDN [53](#)
 - mismatch [53](#)
- host name mismatch
 - change format to FQDN [53](#)

I

- identify
 - available databases [21](#)
 - available resources [21](#)
- information
 - how to send feedback about improving documentation [73](#)
- issues
 - troubleshooting using log files [52](#)

J

jobs

- monitoring progress of restore operations [44](#)
- monitoring progress of SnapCenter clone operations from the Jobs page [51](#)
- monitoring progress of SnapCenter verification operations from the Jobs page [34](#)

Jobs page

- monitoring progress of SnapCenter backup operations [30](#)
- monitoring progress of SnapCenter clone operations [51](#)
- monitoring progress of SnapCenter restore operations [44](#)
- monitoring progress of SnapCenter verification operations [34](#)

jobs, SnapCenter backup

- monitoring progress from the Jobs page [30](#)

jobs, SnapCenter clone

- monitoring progress from the Jobs page [51](#)

jobs, SnapCenter restore

- monitoring progress from the Jobs page [44](#)

jobs, SnapCenter verification

- monitoring progress from the Jobs page [34](#)

L

Linux

- commands for backing up Oracle databases [67](#)
- commands for cloning Oracle databases [69](#)
- commands for restoring and recovering Oracle databases [68](#)

Linux commands

- common tasks for scripting backup, restore, recovery, and clone operations [67](#)

log files

- accessing to troubleshoot issues [52](#)

log messages

- display incorrect time zone [59](#)

logging in

- as a SnapCenter administrator [15](#)
- as a SnapCenter user [15](#)
- to SnapCenter [15](#)
- with more than one role [15](#)

M

managing backups

- delete [65](#)
- rename [65](#)

missing

- database name in oratab file [52](#)
- disk paths [54](#)
- SID in oratab file [52](#)

modifying

- datasets [61](#)
- policies [63](#)

monitoring

- progress of SnapCenter backup operations [30](#)
- progress of SnapCenter clone operations from the Jobs page [51](#)

- progress of SnapCenter restore operations [44](#)
- progress of SnapCenter verification operations from the Jobs page [34](#)

mounting

- database backups [36](#)

N

NFS mount point

- busy [57](#)

non-multipath environment

- data protection fails [59](#)

O

offline backups

- described [18](#)

on-demand backup operations

- from a dataset [30](#)
- from inventory page [29](#)

on-demand backups

- verifying [33](#)

online backups

- described [18](#)

operation, data protection

- fails [59](#)

Operational lock file

- not deleted [58](#)

operations are not executed

- insufficient space [58](#)

operations fail

- insufficient space to create Snapshot copies [58](#)

operations on datasets

- stopping temporarily [61](#)

operations, restore and recovery

- workflow for Oracle databases [38](#)

operations, SnapCenter backup

- monitoring progress from the Jobs page [30](#)

operations, SnapCenter clone

- monitoring progress from the Jobs page [51](#)

operations, SnapCenter restore

- monitoring progress from the Jobs page [44](#)

operations, SnapCenter verification

- monitoring progress from the Jobs page [34](#)

Oracle databases

- clone strategy [46](#)
- creating backup datasets for [27](#)
- creating backup policies for [23](#)
- creating verification policies for [22](#), [32](#)
- identifying resources for backup [21](#)
- Linux commands for backing up [67](#)
- Linux commands for cloning [69](#)
- Linux commands for restoring databases [68](#)
- supported configurations for backup [17](#)
- verify workflow [32](#)
- when to back up [19](#)
- workflow for backing up [16](#)
- workflow for cloning backups [46](#)
- workflow for restoring and recovering [38](#)

overview

- SnapCenter Plug-in for Oracle Database [6](#)

P

- Plug-in for Oracle
 - features [6](#)
 - overview of [6](#)
 - tasks you can perform using [6](#)
- Plug-in for Oracle database
 - backups supported for cloning [47](#)
- Plug-in for Oracle Database
 - prerequisites for using [13](#)
- policies
 - attaching to datasets for Oracle databases [27](#)
 - copying [63](#)
 - deleting [64](#)
 - detaching from a dataset [64](#)
 - modifying [63](#)
 - used in SnapCenter data protection [10](#)
 - viewing [64](#)
- policies, backup
 - creating for Oracle databases [23](#)
 - managing [63](#)
- policies, verification
 - creating for Oracle databases [22](#), [32](#)
 - managing [63](#)
- postscripts
 - supported [21](#)
- preferred nodes
 - RAC setup [19](#)
- prerequisites
 - for using Plug-in for Oracle Database [13](#)
- prescripts
 - supported [21](#)

R

- RAC setup
 - preferred nodes [19](#)
- recovery
 - types of [41](#)
- recovery fails
 - cloned database [56](#)
- recovery operations
 - workflow for Oracle databases [38](#)
- recovery types
 - all logs [41](#)
 - date and time [41](#)
 - SCN [41](#)
- registering backup
 - fails [53](#)
- release notes
 - overview of [6](#)
- renaming
 - backup copies [65](#)
- replication
 - using SnapMirror [20](#)
 - using SnapVault [20](#)
- repositories
 - for SnapCenter, overview of [7](#)
- requirements
 - for using Plug-in for Oracle Database [13](#)
- resources
 - backing up on demand [29](#)
 - used in SnapCenter data protection [10](#)

- resources backup
 - retention options [20](#)
 - when to perform [19](#)
- restore
 - types of [40](#)
- restore and recovery
 - defining a strategy [38](#)
- restore and recovery operations
 - limitations [40](#)
 - using Linux commands for Oracle databases [68](#)
- restore failure
 - datafiles and control files [55](#)
 - from secondary volume [55](#)
 - load-sharing mirror configured [55](#)
- restore methods
 - connect-and-copy [40](#)
 - in-place [40](#)
 - types of [40](#)
- restore operations
 - monitoring [44](#)
 - possible sources and destinations [41](#)
 - resources, datasets, and policies overview [10](#)
 - workflow for Oracle databases [38](#)
- restore operations, SnapCenter
 - monitoring progress from the Jobs page [44](#)
- restore types
 - full [40](#)
 - partial [40](#)
- restoring and recovering
 - databases [41](#)
- restoring and recovering resources
 - strategy [38](#)
- restoring databases
 - destinations for a restore operation [41](#)
 - preparing for [13](#)
 - sources for a restore operation [41](#)
- root file system
 - insufficient space [58](#)

S

- scheduling
 - backup verification [33](#)
- scripts
 - for data protection jobs [21](#)
- secondary volume
 - load-sharing mirror
 - configured on source volume [55](#)
 - SnapMirror and SnapVault configured [55](#)
- SnapCenter
 - components [7](#)
 - finding the URL [15](#)
- SnapCenter backup operations
 - monitoring progress from the Jobs page [30](#)
- SnapCenter clone operations
 - monitoring progress from the Jobs page [51](#)
- SnapCenter Plug-in for Microsoft SQL Server
 - overview of components [7](#)
- SnapCenter Plug-in for Microsoft Windows
 - overview of components [7](#)
- SnapCenter Plug-in for Oracle database
 - overview of components [7](#)
- SnapCenter Plug-in for Oracle Database

- features [6](#)
- overview of [6](#)
- tasks you can perform using [6](#)
- SnapCenter Plug-in for UNIX
 - overview of components [7](#)
- SnapCenter Plug-ins for Linux operation
 - fails [59](#)
- SnapCenter restore operations
 - monitoring progress from the Jobs page [44](#)
- SnapCenter Server
 - overview [7](#)
- SnapCenter verification operations
 - monitoring progress from the Jobs page [34](#)
- SnapMirror
 - use of [20](#)
- Snapshot copies
 - insufficient space to create [58](#)
- SnapVault
 - use of [20](#)
- storage discovery process
 - back up operation fails [53](#)
- storage types
 - supported by plug-ins package for Linux [14](#)
- suggestions
 - how to send feedback about documentation [73](#)
- SVMs
 - supported storage types [14](#)

T

- terminology
 - resources, datasets, and policies data protection [10](#)
- troubleshooting
 - accessing log files to [52](#)
 - backup failure [52](#)
 - data protection operations [52](#)
- Twitter
 - how to receive automatic notification of documentation changes [73](#)

U

- unmounting

- database backups [36, 37](#)
- users
 - using roles at login [15](#)

V

- verification failure
 - datafiles backup [54](#)
- verification operations, SnapCenter
 - monitoring progress from the Jobs page [34](#)
- verification policies
 - creating for Oracle databases [22, 32](#)
 - managing [63](#)
- verify operations
 - workflow for Oracle databases [32](#)
- verifying
 - backups on demand [33](#)
- verifying backups
 - on primary or secondary storage [21](#)
 - on SnapMirror or SnapVault storage [21](#)
 - scheduling [20](#)
- view
 - clone details [66](#)
- viewing
 - policies [64](#)
- VMDK backup
 - fails [52](#)
- VMs
 - supported storage types [14](#)

W

- web address
 - for SnapCenter [15](#)
- workflows
 - data protection diagram [12](#)
 - for backing up Oracle databases [16](#)
 - for cloning Oracle database backups [46](#)
 - for restoring and recovering Oracle databases [38](#)
 - verifying Oracle databases [32](#)