



Technical Report

SnapCenter Plug-In for Oracle Database

Best Practices

Ebin Varghese Kadavy, Jeffrey Steiner, Subhash Athri, NetApp
June 2018 | TR-4700

Abstract

NetApp® SnapCenter® is a unified, scalable platform for Oracle-consistent data protection that automates complex operations with centralized control and oversight. This report explains the best practices for deploying and ensuring data availability of Oracle database deployments on NetApp ONTAP® storage running in an on-premises, colocated data center and cloud.

TABLE OF CONTENTS

1	Introduction	4
1.1	Audience	4
1.2	Purpose	4
2	SnapCenter Overview	4
2.1	Architecture	4
2.2	SnapCenter Features	7
3	SnapCenter Best Practices for Oracle Database	7
3.1	Best Practices for Storage Volume/LUN Layout	7
3.2	Best Practices for Oracle ASM Configurations	8
3.3	Supported and Unsupported Features and Layouts	8
4	Preinstallation Best Practices for SnapCenter	9
4.1	SnapCenter Server Resource Requirements	9
4.2	SnapCenter Server-Package, Storage, and Browser Requirements	9
4.3	Connection and Port Requirements	10
4.4	SnapCenter Plug-In for Oracle Host Requirements	11
5	Oracle Database Plug-In Configuration Best Practices	12
6	SnapCenter Policies and Resource Group Best Practices	13
7	Oracle Database Backup Best Practices	16
8	Restore and Recovery Best Practices	18
8.1	Archive Log Management for Advanced Recovery	20
8.2	Restore from Secondary Mirror or Vault Storage	23
8.3	Disaster Recovery	26
8.4	Important Considerations	27
9	Oracle Clone Best Practices	27
9.1	To Perform a Clone Operation by Using the Clone Wizard	29
9.2	To Perform a Clone Operation by Using the CLI	32
9.3	RAC-to-RAC Clone	32
9.4	Important Considerations	33
	Appendix	34
A.	SnapCenter Deployment Models for Oracle Database	34
B.	RAC One Node and Other Third-Party Cluster Solutions (Active-Passive)	36
C.	Block-Level Recovery	37

D. Tablespace Point-in-Time Recovery	39
E. Recover a Table.....	39
Where to Find Additional Information	42
Version History	42

LIST OF TABLES

Table 1) SnapCenter Server resource requirements (specific to Oracle deployments).....	9
Table 2) SnapCenter server-package, storage, and browser requirements.	9
Table 3) Connection and port requirements.	10
Table 4) SnapCenter Plug-In for Oracle host requirements.....	11

LIST OF FIGURES

Figure 1) SnapCenter architecture	5
Figure 2) SnapCenter scaled-out deployment.....	6
Figure 3) Resource group creation.....	15
Figure 4) SnapCenter restore.....	19
Figure 5) SnapCenter clone	28
Figure 6) Private data center/cloud deployment.....	34
Figure 7) NetApp Private Storage deployment.....	35
Figure 8) Hybrid cloud deployment.....	35

1 Introduction

In today's data-driven enterprise, business-critical applications must be operational around the clock to facilitate decision making, e-commerce, and many other business processes. Virtualization has improved significantly and in recent years has gradually become the platform of choice for tier 1 applications. Customers who use enterprise applications like Oracle, Microsoft SQL Server, Exchange, and SAP are starting to choose virtualization over physical deployments.

Oracle databases are among the most important applications in many environments. Relational databases are used for custom applications developed internally by a company, or as a database back end for commercial application deployments. In both scenarios, data registered in Oracle databases requires proper design to be readily available. Rapid growth in data volume and application demands make it increasingly difficult to provide availability and protection for valuable data assets.

Administrators need tools that enable them to take frequent backups with minimal impact on operations, perform quick application recovery, and rapidly create copies for user testing and development regardless of physical, virtual or hybrid cloud deployments. SnapCenter data protection software offers a high degree of availability for Oracle databases by leveraging its capabilities such as data loss avoidance, verified protection, and high-speed recovery.

1.1 Audience

This document is intended for use by customers who are using Oracle databases in both physical and virtualized environments. It is a source of useful information about best practices on data protection for Oracle databases for storage administrators, database administrators, virtualization specialists, architects, and data protection administrators.

1.2 Purpose

This document describes the best practices for deploying and ensuring data availability of Oracle deployments on NetApp ONTAP® storage running on-premises, in a colocated data center or in the cloud and leveraging SnapCenter Plug-In for Oracle. The recommendations in this report are generic; they are not specific to any configuration and, depending on your business needs, some suggestions might require changes. Each environment should be carefully evaluated against the official documentation for SnapCenter, hypervisor vendors, and Oracle.

2 SnapCenter Overview

SnapCenter is NetApp's next-generation data protection software for tier 1 and tier 2 enterprise applications. With its "single pane of glass" management interface, SnapCenter automates and simplifies the manual, complex, and time-consuming processes associated with the backup, recovery, and cloning of multiple databases and other application workloads.

SnapCenter leverages NetApp technologies, including Snapshot™ copies, SnapMirror® replication software, SnapRestore® data recovery software, and FlexClone® volumes, that allow it to integrate seamlessly with technologies offered by Oracle, Microsoft, SAP, VMware, and Mongo across FC, FCoE, iSCSI, and NAS protocols. This integration allows IT organizations to scale their storage infrastructure, meet increasingly stringent SLA commitments, and improve the productivity of administrators across the enterprise.

2.1 Architecture

Figure 1 shows the SnapCenter architecture.

Figure 1) SnapCenter architecture

Simple Deployment

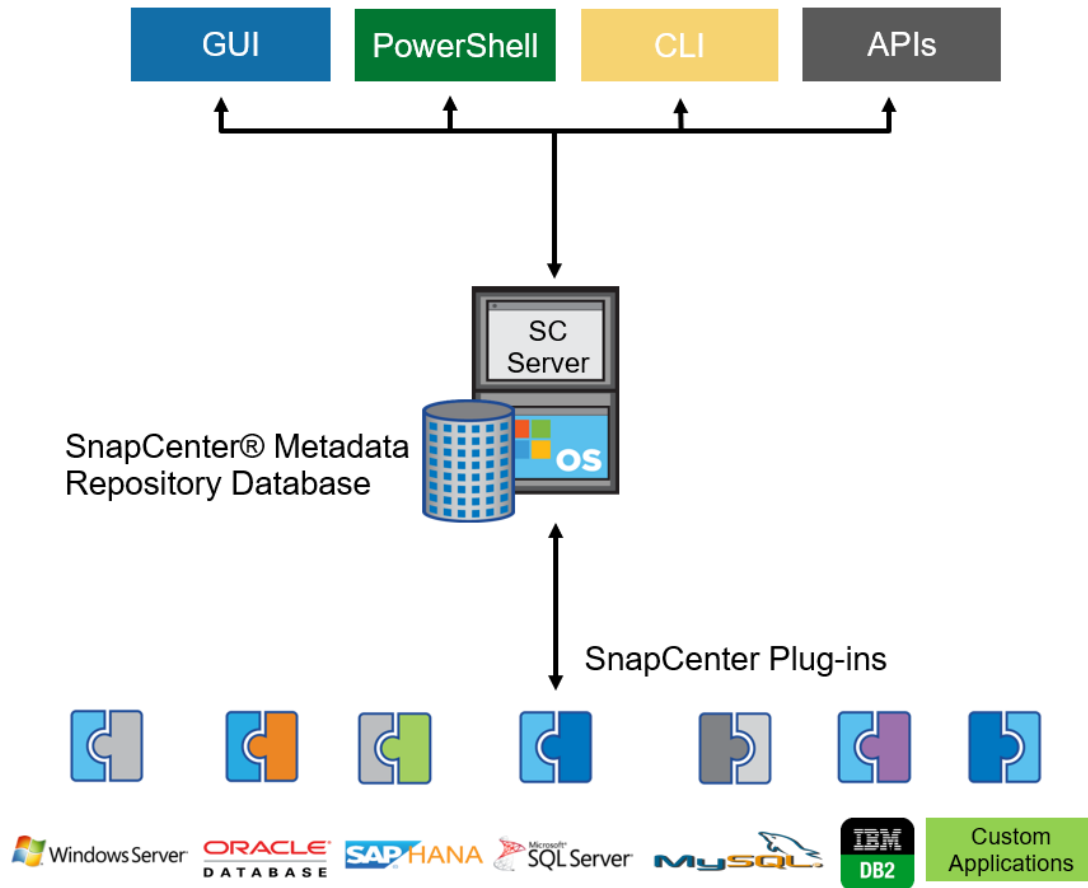


Figure 2) SnapCenter scaled-out deployment

Scaled-Out Deployment

VMware Integration

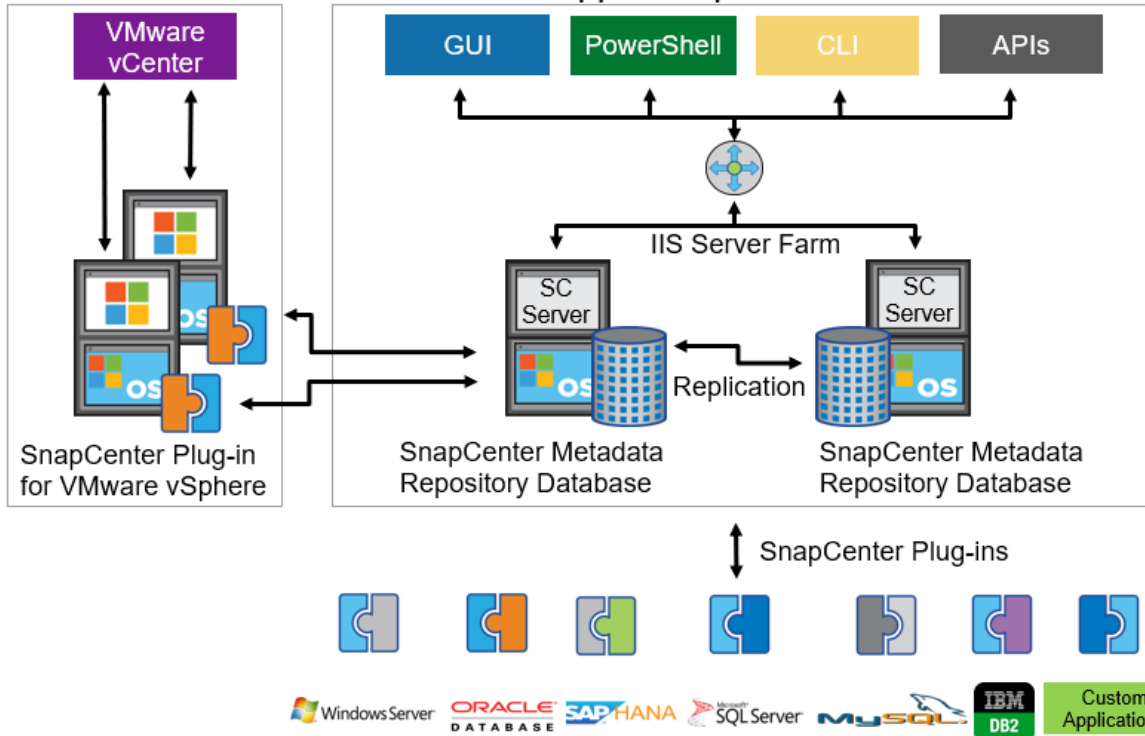


Figure 2 shows a scaled deployment of SnapCenter that gives maximum availability at both the compute and repository layers.

SnapCenter comes with four user interfaces: CLI, GUI, PowerShell, and REST. The SnapCenter metadata repository database (MySQL), which holds the backup, restore, and clone metadata, can be scaled up to two nodes to support high availability (HA). HA and load balancing for the SnapCenter Server are handled by the Microsoft Network Load Balancer (NLB) and Application Request Routing (ARR).

SnapCenter 4.0 offers plug-ins for Oracle database, Microsoft SQL Server, SAP HANA, Exchange, MySQL, Mongo DB, IBM DB2, VMware virtual machines, Windows file system, and storage.

SnapCenter Plug-In for Oracle has two major components:

- **SMCore** is a core component of SnapCenter that runs colocated within the SnapCenter server host, unlike other plug-ins. It plays a significant role in invoking and handling all the major workflows, starting with discovering the storage devices, hosting the file system, secondary replication of backups, and retention management. SMCore coordinates with the Linux Plug-In (managed by the SnapCenter plug-in loader (SPL)) to perform Oracle database workflows.
- **SPL** is a component that runs on the Oracle database host that loads and runs the Oracle plug-in. SMCore coordinates with SPL to perform Oracle data protection workflows like quiesce/unquiesce, backup, RMAN catalog, mount, restore, and clone.

2.2 SnapCenter Features

SnapCenter provides a centralized management environment, while using role-based access control (RBAC) to delegate data protection and management capabilities to individual application users across the SnapCenter Server and Windows or Linux hosts. SnapCenter includes the following key features:

- A unified and scalable platform across applications and database environments and virtual and nonvirtual storage, powered by the SnapCenter Server
- Consistency of features and procedures across plug-ins and environments, supported by the SnapCenter user interface
- RBAC for security and centralized role delegation
- Application-consistent Snapshot management, restore, clone, and backup verification support from both primary and secondary destinations (NetApp SnapMirror and SnapVault® backup software)
- Remote package installation from the SnapCenter GUI
- Nondisruptive, remote upgrades
- A dedicated SnapCenter repository for fast data retrieval
- Load balancing implemented by using NLB and ARR, with support for horizontal scaling
- Centralized scheduling and policy management to support backup and clone operations
- Centralized reporting, monitoring, and dashboard views

3 SnapCenter Best Practices for Oracle Database

NetApp recommends the following best practices for the Oracle database.

3.1 Best Practices for Storage Volume/LUN Layout

SnapCenter does not require having dedicated volume/LUN layouts for provisioning the Oracle database. However, it is a good practice to isolate the Oracle data files from all other Oracle files such as control file, archive log and redo log file, spfile/pfile, password file, and so on. This isolation enables faster recovery (using the ONTAP SnapRestore mechanism) and addresses disaster recovery scenarios. For more information about the volume layouts, refer to [Oracle Databases on ONTAP TR-3633](#).

Best Practices for Dedicated storage (Volume/LUN) Layout

- Place Oracle data files in the dedicated volumes/LUNs with no other files (like physical copies of data files, backups, scripts, or plain text files, including Oracle files). This best practice holds good for physical and virtualized VMware layouts however VMFS (VMware Virtual Machine File System) layouts are considered to be an exception as it makes use of VMware's storage VMotion rather than using storage SnapRestore.
- For SAN layouts, especially in large database environments with Oracle Automatic Storage Management (ASM), NetApp recommends keeping the disk group for the Oracle data files (DATA) in the separate volumes where LUNs or ASM disks are dedicated only for data files and not for any other files.

Best Practices for Shared Storage (Volume/LUN) Layout

- Shared volume/LUN layouts are a better fit for smaller databases and noncritical databases. It might be possible to cut down the overall operational and management cost by sharing multiple databases in the same volume.
- When multiple databases are hosted in the same volume/LUN, and if these databases reside on the same host and are backed up together in a resource group, then Snapshot consolidation optimizes the system to take a single Snapshot copy for all the databases running on the same volume.

- The advantage of NFS over VMFS is that the NFS datastore considers each VMDK as a single file, so that data files deployed on separate VMDKs can be quickly restored by using ONTAP Single File SnapRestore (SFSR) technology.
- You can also share the archive log destination for all the databases running in a host within the same volume/LUN as the Oracle flash recovery area (FRA) or non-FRA destination. During backup when multiple databases are grouped together in a resource group, one Snapshot copy is taken in the storage for the entire archive log volume for all the databases thus consolidating everything in a single Snapshot.
- When multiple databases share same volume/LUN, the retention of backups might be affected. For example: If 100 databases share the same volume and if daily Snapshot copies are triggered for all 100 databases then you cannot retain the Snapshot copies for more than 2.5 days because you will reach the 255 Snapshot limit very easily. Hence it is advised to keep a limited number of databases sharing the same volume/LUN to avoid retention bottlenecks.

3.2 Best Practices for Oracle ASM Configurations

NetApp recommends the following best practices for Oracle ASM configurations:

- For ASM layouts, make sure that the `ASM_DISKSTRING` value is set appropriately before handling any SnapCenter operations. If this value is not set correctly, then restore, clone, or mount operations fail.
 - If you are using the ASMLIB package for managing ASM devices, set the `ASM_DISKSTRING` value to `ORCL:*`.
 - For non ASMLIB scenarios such as `udev` rules, set the value to `/dev/< exact device location >`.
- SnapCenter doesn't support ASM disks that have multiple partitions. It can only support disks with single or no partitions.
- If ASM with multipathing on Linux is used, make sure that `/etc/sysconfig/oracleasm` has the following variables set:

```
ORACLEASM_SCANORDER='dm'
ORACLEASM_SCANEXCLUDE='sd'
```

3.3 Supported and Unsupported Features and Layouts

SnapCenter supports:

- RAC, ASM (over NFS, physical SAN), Data Guard Standby, and Active Data Guard Standby
- NFS v3, SAN ext3, ext4, XFS file system
- iSCSI, FC, FCoE
- Linux LVM, file system on RAW device, NFS and VMFS datastore layouts, RDM
- Backup, clone, and restore of Data Guard and Active Data Guard
- RAC One Node and third-party cluster solutions (active-passive)

SnapCenter does not support:

- ASM over VMDKs in NFS/VMFS datastores and RDM
- Disaster recovery situations; in case of failure of the entire site, the SnapCenter Server repository cannot be recovered to a different site
- Golden Gate, Databroker, ASMFD, ACFS
- Third-party volume manager solutions
- Point-in-time recovery of tablespace and pluggable database (PDB) data file restore
- Volume-based SnapRestore

- Encrypted file system, autofs, LSM
- Multiple SAN igroups for a single LUN
- Recovery of Data Guard and Active Data Guard
- Backup verification on remote host

Note: FlexASM and FlexCluster are new features in Oracle 12c. All SnapCenter workflows should go through until they encounter an ASM failover scenario. This is a known limitation and it will be addressed in subsequent releases.

4 Preinstallation Best Practices for SnapCenter

Preinstallation best practices cover requirements and prerequisites for both SnapCenter Server and the Oracle plug-in. Here are the minimum requirements and prerequisites to consider before installing a SnapCenter Server for managing Oracle environments.

4.1 SnapCenter Server Resource Requirements

Table 1) SnapCenter Server resource requirements (specific to Oracle deployments).

Item	Smaller Environments (<30 DBs) Mission- and Business-Critical Oracle Databases	Larger Environments (<100 DBs) Mix of Critical & Dev/Test Oracle Databases	Very Large Environments (<300 DBs) Mix of Critical Production and Dev/Test Oracle Databases
RAM (Based on frequency of jobs on each host)	8-16G	32G	64G
CPUs	4	8	8
Hard-drive space for SnapCenter software and logs Minimum of 3 years (Based on retention of jobs or log pruning)	10-20GB	30GB	> 50GB
Hard-drive space for SnapCenter repository (metadata) Minimum of 3 years (Based on the frequency of retaining backups)	20-30GB	60GB	100GB

4.2 SnapCenter Server-Package, Storage, and Browser Requirements

Table 2) SnapCenter server-package, storage, and browser requirements.

Item	Requirements
Operating System	Windows 2012 R2, 2016

Item	Requirements
Packages	<p>Standalone single-instance server deployment:</p> <ul style="list-style-type: none"> • Microsoft .NET Framework 4.5.2 or later • Windows Management Framework (WMF) 4.0 or later • PowerShell 4.0 or later • Java 1.8 (64-bit) <p>Cluster or high-availability deployment:</p> <ul style="list-style-type: none"> • IIS • Web Platform Installer 3.0 • Web Platform Installer 5 (upgrade from 3.0) • Web Deploy 2.0 • Web Deploy 3.5 (upgrade from 2.0) • Web Farm Framework 2.0 (requires IIS and Web Platform Installer 3.0) • Application Request Routing 3.0 • URL Rewrite Module 2.0 (installed by ARR 3.0) • External Disk Cache (installed by ARR 3.0) <p>For the latest information about supported versions, see the NetApp Interoperability Matrix Tool (IMT).</p>
Storage requirement	<p>Works with ONTAP 8.2.2 and later</p> <p>NetApp recommends running applications and VMs on ONTAP 9.x for newer features, better scalability, and higher performance. Consider upgrading earlier ONTAP systems to 9.x to take advantage of these benefits.</p>
Browser requirements	<p>Internet Explorer, Edge, Google Chrome</p> <p>For the latest information about supported versions, see the NetApp Interoperability Matrix Tool (IMT).</p>

4.3 Connection and Port Requirements

SnapCenter uses various ports for each plug-in; some ports can be customized and others cannot. Table 3 lists the ports that must be enabled or freed from a firewall before installing SnapCenter software.

Table 3) Connection and port requirements.

Type of Port	Default Port Number and Purpose
SnapCenter port (server)	<p>8146 (HTTPS), bidirectional, customizable</p> <p>SnapCenter URL: https://<ip address/hostname>:8146</p> <p>This port cannot be changed post-installation.</p>
SnapCenter SMCore communication port (server)	<p>8145 (HTTPS), bidirectional, customizable</p> <p>This port is used for communication between the SnapCenter Server and the hosts where the SnapCenter plug-ins are installed.</p>
MySQL port (server)	<p>3306 (HTTPS), bidirectional, fixed</p> <p>This port is used for communication between SnapCenter and the MySQL repository database.</p>

Type of Port	Default Port Number and Purpose
Linux plug-in hosts	<p>22 (SSH), not customizable</p> <p>These ports are used by SnapCenter to copy plug-in package binaries to Linux plug-in hosts. They should be open or excluded from the firewall or iptables.</p> <p>8145 (HTTPS), bidirectional, customizable</p> <p>This port is used for communication between SMCORE and hosts where the SnapCenter plug-ins are installed. The communication path must be open between the SVM management LIF and the SnapCenter Server.</p>
SnapCenter Plug-In for VMware vSphere port	<p>8144 (HTTPS), bidirectional, customizable</p> <p>Used for communications from the vCenter vSphere web client and from the SnapCenter Server.</p> <p>Note: You cannot modify the port if the plug-in is installed on the SnapCenter Server host.</p>
VMware vSphere vCenter Server port	<p>443 (HTTPS), bidirectional</p> <p>The port is used for communication between the host for the SnapCenter Plug-In for VMware vSphere and vCenter.</p>

SnapCenter expects its Linux host to be resolvable during the host registration process.

If the Linux host or storage cannot be resolved to a fully qualified domain name (FQDN) for any reason (for example, the host coming from a different cloud, different domain, or private network), the workaround is to add an entry of the host or storage with a hostname (FQDN) in the

C:\Windows\System32\drivers\etc\hosts file in Windows.

If you are adding a host running in any public cloud and you plan to manage databases (backup, restore, and clone) by using SnapCenter running in either a private data center or cloud, it is important to make sure that the firewall is open to listen to host IPs and ports 8145 and 8146 with inbound and outbound communication enabled. For example, if you want to clone an Oracle database running on premises to Amazon Web Services (AWS) cloud, you must enable ports 8145 and 8146 bidirectionally, along with port 22 (SSH) to install the Linux plug-in on the Linux host running in the cloud. However, you can skip port 22 from the firewall list by manually installing the plug-in directly to the Linux host.

4.4 SnapCenter Plug-In for Oracle Host Requirements

Table 4) SnapCenter Plug-In for Oracle host requirements

Item	Requirements
Operating systems	<p>RHEL, Oracle Linux, SUSE</p> <p>Note: If you are using Oracle database on LVM in Oracle Linux or the Red Hat Enterprise Linux 6.6 or 7.0 operating system, you must install the latest version of Logical Volume Manager (LVM).</p> <p>For the latest information about supported versions, see the NetApp Interoperability Matrix Tool (IMT).</p>
Minimum memory for plug-in	<p>Minimum 512MB</p> <p>Recommended 2GB</p>

Item	Requirements
Minimum installation and operational hard-drive space (for a minimum of 3 years)	<p>Default installation location (binaries): <code>/opt/NetApp/SnapCenter</code></p> <ul style="list-style-type: none"> Installation location can be customized (can also be on NetApp storage) 3GB recommended <p>Logs and config file location: <code>/var/opt/SnapCenter/</code></p> <ul style="list-style-type: none"> Log location cannot be changed. 30GB required (3 years, if logs are not pruned repeatedly)
Required software packages	<p>Java 1.8.x (64-bit) Oracle Java and OpenJDK</p> <p>Note: If you are using multiple versions of Java for different applications, you must verify that the <code>JAVA_HOME</code> option is set to the correct Java path for SnapCenter.</p>
Oracle Plug-In Installation user	<p>Default Linux user is root</p> <p>Note: Nonroot accounts work, but sudo privileges are required for the nonroot account. Make sure that the sudo accounts have Oracle database group privilege.</p> <p>NetApp recommends using a sudo package version 1.8.7 or later for checksums.</p> <p>For more information on setting sudo privileges for executables, see the SnapCenter 4.0 Installation and Setup Guide.</p>

5 Oracle Database Plug-In Configuration Best Practices

When registering a new Oracle database host with SnapCenter, keep the following best practices in mind:

- If the hostip can't reverse resolve to a proper FQDN but can be reached from the SnapCenter Server by using its IP address, then you must have an entry of the hostname with FQDN along with the IP address in the `/etc/hosts` file on the SnapCenter Server to complete the host plug-in installation.
Note: Confirm that the `/etc/oratab` file has the correct Oracle home entries for each database. SnapCenter discovers the details of the database only from this location.
- Once the agent is up and running in the Oracle host, the rest of the communication from server to host or vice versa is only through REST API calls and not through the OS account.
- If SSH port 22 can't be granted in the firewall rules, then perform a manual installation of the SnapCenter for Oracle Database Plug-In on the given Linux host. After manually installing the plug-in on the host, you must register the host with SnapCenter. While registering, select the Skip Prechecks checkbox to register the Oracle host with SnapCenter.
- For SAN, `sg3_utils` is required. If LVM is used then the LVM utils package must be installed.
- To enable a port on the Linux host firewall settings such as RHEL 7.1, use the following commands:

```
root> /sbin/iptables -A INPUT -p tcp -m tcp --dport 8145 -j ACCEPT
root> /sbin/iptables -A OUTPUT -p tcp -m tcp --dport 8145 -j ACCEPT
root> service iptables save
```

- If the Oracle database is running on purely VMware virtualized layouts (VMDK or RDM) then you should configure the SnapCenter for VMware vSphere (SCV) plug-in. This plug-in can be installed on

the same Windows Server where the SnapCenter server is deployed or on a different Windows machine. Verify that port 8144 is open in the firewall.

- If databases reside on a VMDK file system, you must log into vCenter and navigate to **VM OPTIONS > ADVANCED > EDIT CONFIGURATION** to set the value of `disk.enableUUID` to `true` for the virtual machine. The VM needs to be rebooted after the change.
- You can use the CLI to automate the installation of the plug-in to the host:
 - a. Perform silent installation of the plug-in with this command:

```
SnapCenter_linux_host_plugin.bin -i silent -DPORT=8145 -DSERVER_IP=xxxx-DSERVER_PORT=8146
```

- b. Run the `add-smHost` cmdlet in PowerShell to register the Linux host with the SnapCenter Server. Refer to the [Cmdlet Reference Guide for Windows](#) for the syntax.
- SnapCenter doesn't allow you to back up both the VMs and the Oracle application together. VMs must be backed up separately with different policies, using the vSphere web client where a lightweight SCV plug-in is integrated. The Oracle database must be backed up from the SnapCenter GUI or the plug-in CLI, or by using SnapCenter PowerShell cmdlets.
 - If you have a RAC or RAC One Node database, select the checkbox to add the cluster nodes in ADD HOST for RAC AWARENESS and verify that the plug-ins are pushed on all the cluster nodes or manually installed on them.
 - Once the plug-in is installed and configured on the Oracle host, the databases should appear in the resources screen. If it's a RAC database, you can enable preferred nodes for backup in the Configure Database wizard. This serves two purposes: to handle host/instance failures and to dedicate the backups to a separate node to avoid load across the cluster.

Additional Resources

For more information on how to register an Oracle database host with SnapCenter, see the following [YouTube demonstration video](#).

To ensure if the Oracle host is successfully configured, check the following [YouTube demonstration video](#)

6 SnapCenter Policies and Resource Group Best Practices

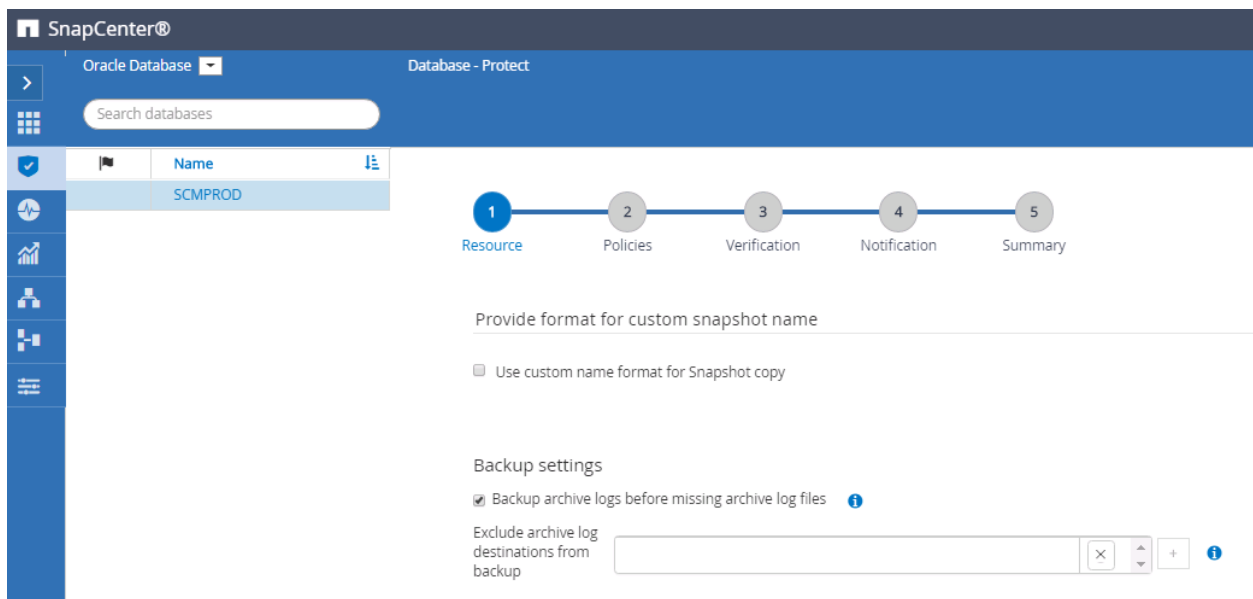
Backup policies are designed based on your business requirements, so it is important to be aware of these essential tips before configuring them:

- Determine if you need a full or partial backup of databases, and whether backups should be online or offline:
 - SnapCenter supports both online and offline backups. The full backup captures data files, archive logs, and control files. SnapCenter does not back up pfile/spfile, password files, listeners, tnsname files, or redo logs (including Oracle homes).
 - SnapCenter also allows you to back up, restore, and clone Oracle Data Guard and Active Data Guard standby databases. To back up a Data Guard database (RAC or non RAC), the policy should have the backup type set to `offline mount backup`. For an Active Data Guard database, the backup type should be `online`.
 - SnapCenter supports using the offline shutdown/mount policy to back up databases without enabling archive logs.
- Define the recovery point objectives (RPO) for your production and nonproduction systems:
 - In SnapCenter terms, RPO can be identified as the backup frequency, that is how frequently you schedule the backup to reduce data loss. The best practice is to schedule archive log backups more frequently than data backups or full backups. The minimum interval can be as short as 15 minutes.

- Snapshot copies can be taken either from SnapCenter or directly from ONTAP storage. The Snapshot copies taken directly from ONTAP are crash consistent; the Snapshot copies taken by SnapCenter are application consistent. Enabling both options at frequent intervals might put you at the risk of hitting the limit of 255 Snapshot copies per volume.
- Define the retention and replication requirements:
 - How long do you want to retain these backups, based on the frequency of backups (hourly, daily, weekly, monthly)?
 - You can retain 255 Snapshot copies in the primary storage and 251 in the secondary disaster recovery or vault storage. You can set retention separately for both on-demand and specific backup frequencies (hourly, daily, weekly, monthly, yearly).
 - To protect the Snapshot copies in the near term, you must replicate your backups to a SnapVault destination. The SnapVault or SnapMirror replication relationship must be done outside of SnapCenter before creating the backup policy.
 - Retention for backups replicated in secondary storage must be handled directly from ONTAP.
- Determine if you want to verify and validate each backup:
 - Verification is an optional feature that can be used to validate the files that are part of the backup. Verification can be enabled in the backup policy and activated during the protection workflow. You can verify the backups from both primary and secondary disaster recovery or vault storage. You can also defer the verification by scheduling it for a later time either from the GUI or the CLI. During verification, a FlexClone volume is created from the backup and mounted to the host, and the Oracle `dbv` utility is run across all the data files present in the newly mounted FlexClone volume.
- Determine if you want to catalog the backups with RMAN for advanced or granular recovery:
 - SnapCenter allows you to catalog the NetApp Snapshot copies with Oracle RMAN. This optional feature helps with granular restores and recoveries such as block-level recovery, tablespace PITR, and so on. During the catalog process, a FlexClone volume is taken at the storage level, the volume/LUNs or disk groups are mounted to the host and the RMAN catalog operation is run across the file system to capture the metadata. Once the metadata is captured, the mounted disk groups, volumes, or LUNs are released or deleted. The following screenshot shows the backup policies.

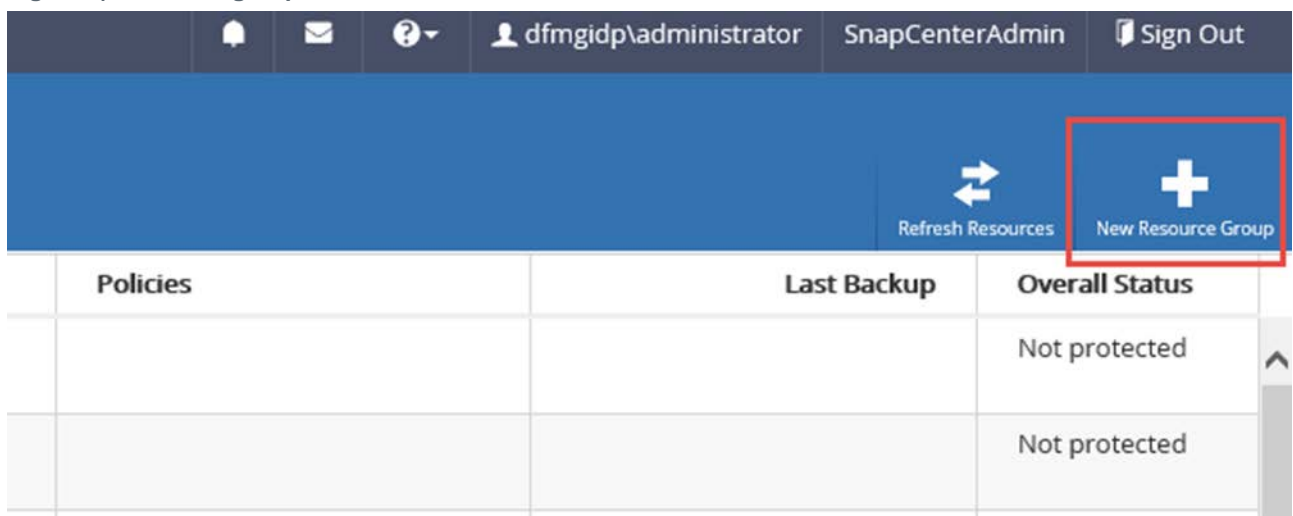
Name	Backup Type	Schedule Type	Replication	Verification
ERP PROD DAILY	FULL, ONLINE	Daily		
ERP_GOLD_BACKUP	FULL, ONLINE	Daily	SnapVault	
PAYB_WITH_DR	FULL, ONLINE	Hourly	SnapMirror	
SCM Daily Gold backup	FULL, ONLINE	Daily	SnapVault	
TEST BACKUP	FULL, ONLINE	Daily	SnapVault	

- Once the backup policies have been created, you must assign them to a database during the protect workflow. The protect workflow internally creates an implicit resource group for that database to enable the backup. You can reuse the same policy for other databases. The following screenshot shows implicit resource group creation for database SCMPROD database protection.



- If you have multiple databases to be backed up in a short window, consider grouping them in a resource group. It's the same as the protect workflow for an individual resource.

Figure 3) Resource group creation.



- During resource group creation, if you group multiple databases in the same host that shares the same volume or LUNs for a backup, you benefit from the Snapshot consolidation feature, which means that one Snapshot copy is taken for all databases instead of multiple Snapshot copies for each database. This also reduces the overall storage overhead and might prevent the system from reaching the Snapshot limit.
- If you have created a test policy and assigned it to multiple databases for backups, you must find the databases that share the same policy and delete the backups first before deleting the entire policy.

7 Oracle Database Backup Best Practices

This section describes the best practices, guidelines, and procedures for Oracle database backup.

To Back Up by Using the CLI

1. To perform a backup using the CLI, use the `sccli` located in `/opt/NetApp/SnapCenter/spl/bin`.
2. Open the connection to the SnapCenter Server running on a Windows Server:

```
[root@rhel3 bin]# ./sccli open-SmConnection
INFO: A connection session will be opened with SnapCenter 'https://SnapCtr.demo.netapp.com:8146/'.
Enter the SnapCenter user name: demo\administrator
Enter the SnapCenter password:

INFO: A connection session with the SnapCenter was established successfully.
```

If you want to keep the connection open until reboot of the server, pass an additional parameter `TokenNeverExpires` to disable token expiry. For security reasons, NetApp does not recommend keeping the token open.

```
[root@rhel3 bin]# ./sccli open-SmConnection -TokenNeverExpires
INFO: A connection session will be opened with SnapCenter 'https://SnapCtr.demo.netapp.com:8146/'.
Enter the SnapCenter user name: demo\administrator
Enter the SnapCenter password:
```

3. After creating backup policies and enabling protection for a resource or resource group, run the following command to back up the entire database (data file and archive log).

```
[root@rhel3 bin]# ./sccli New-SmBackup -policy 'Oracle Daily Online Full' -resource 'host=rhel3,type=Oracle Database,names=[SCMPROD]'

INFO: Job 'Backup of Resource Group 'RHEL3_demo_netapp_com_SCMPROD' with policy 'Oracle Daily Online Full'' QUEUED with jobId '29'
INFO: The command 'New-SmBackup' executed successfully.
[root@rhel3 bin]#
```

4. You can check for additional parameters and syntaxes by using `-help`.

```
./sccli -help
./sccli New-SmBackup -help
```

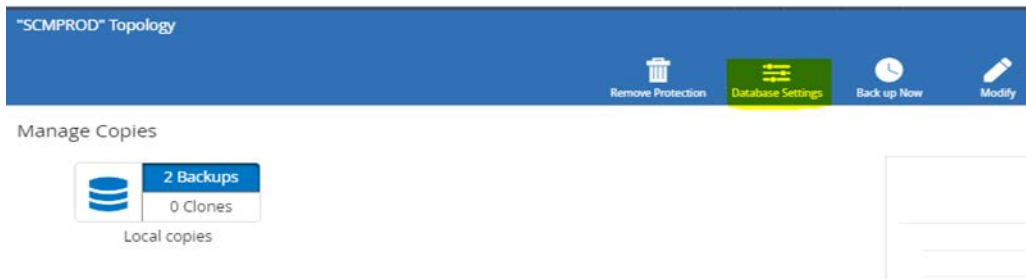
If you have enabled unified replication — that is, both mirror and vault replication — then backups are shown in both the mirror and vault destinations in the topology view.

The retention logic for backups is checked only at the end of each backup job. If any Snapshot (backup) copy is locked by a FlexClone volume (due to clone or mount operations), the retention skips the current Snapshot copy that is locked and moves on to the next one.


SnapCenter backup can be cataloged either on the target control file or the catalog database. If you leave catalog settings unconfigured, then by default SnapCenter catalogs the backups only with the target control file. Follow these steps to configure the catalog database connectivity.

To Configure Catalog Database Connectivity

1. Go to **Topology view of a database > Database Settings > Configure Database**.
2. Click **Database Settings**.



3. Click **Configure Database** to update the settings.



To Configure the RMAN Catalog settings

If RMAN cataloging is enabled, the RMAN tag is automatically generated by SnapCenter to identify the backups.

☺ [Configure RMAN Catalog settings](#) ⓘ

Use Existing Run As +

TNS Name

1. To check the SnapCenter backup details that are cataloged with RMAN, connect to the RMAN prompt and then enter the following commands:

```
Rman> list datafilecopy all
Rman> list copy of archivelog all
```

Note: The data file backups that are cataloged with RMAN are identified by the unique tag `SCO_<DBname>` as a prefix.

2. If archive logs have been deleted outside of SnapCenter or RMAN, enable the parameter `ENABLE_CROSSCHECK=true` in the `sco.properties` file to avoid unexpected delays in searching for the stale archive log entries during backup. The `sco.properties` file is located in `/var/opt/snapcenter/sco/etc`.
3. When a large amount of data is being transferred, SnapVault replication of backups might time out. Therefore you should include the parameters listed below with higher values and specify the values in the `<appSettings>` section of the `SMCoreServiceHost.exe config` file located under `C:\Program Files\NetApp\SMCore` in the SnapCenter Server.

Restart the SnapCenter SMCore service (timeout values are in milliseconds.)

```
<appSettings>
<add key="SnapmirrorRetry" value="288"/>
<add key="SnapmirrorTimeout" value="300000"/>
<add key="SnapshotCheckRetry" value="288" />
<add key="SnapshotCheckTimeout" value="300000" />
</appSettings>
```

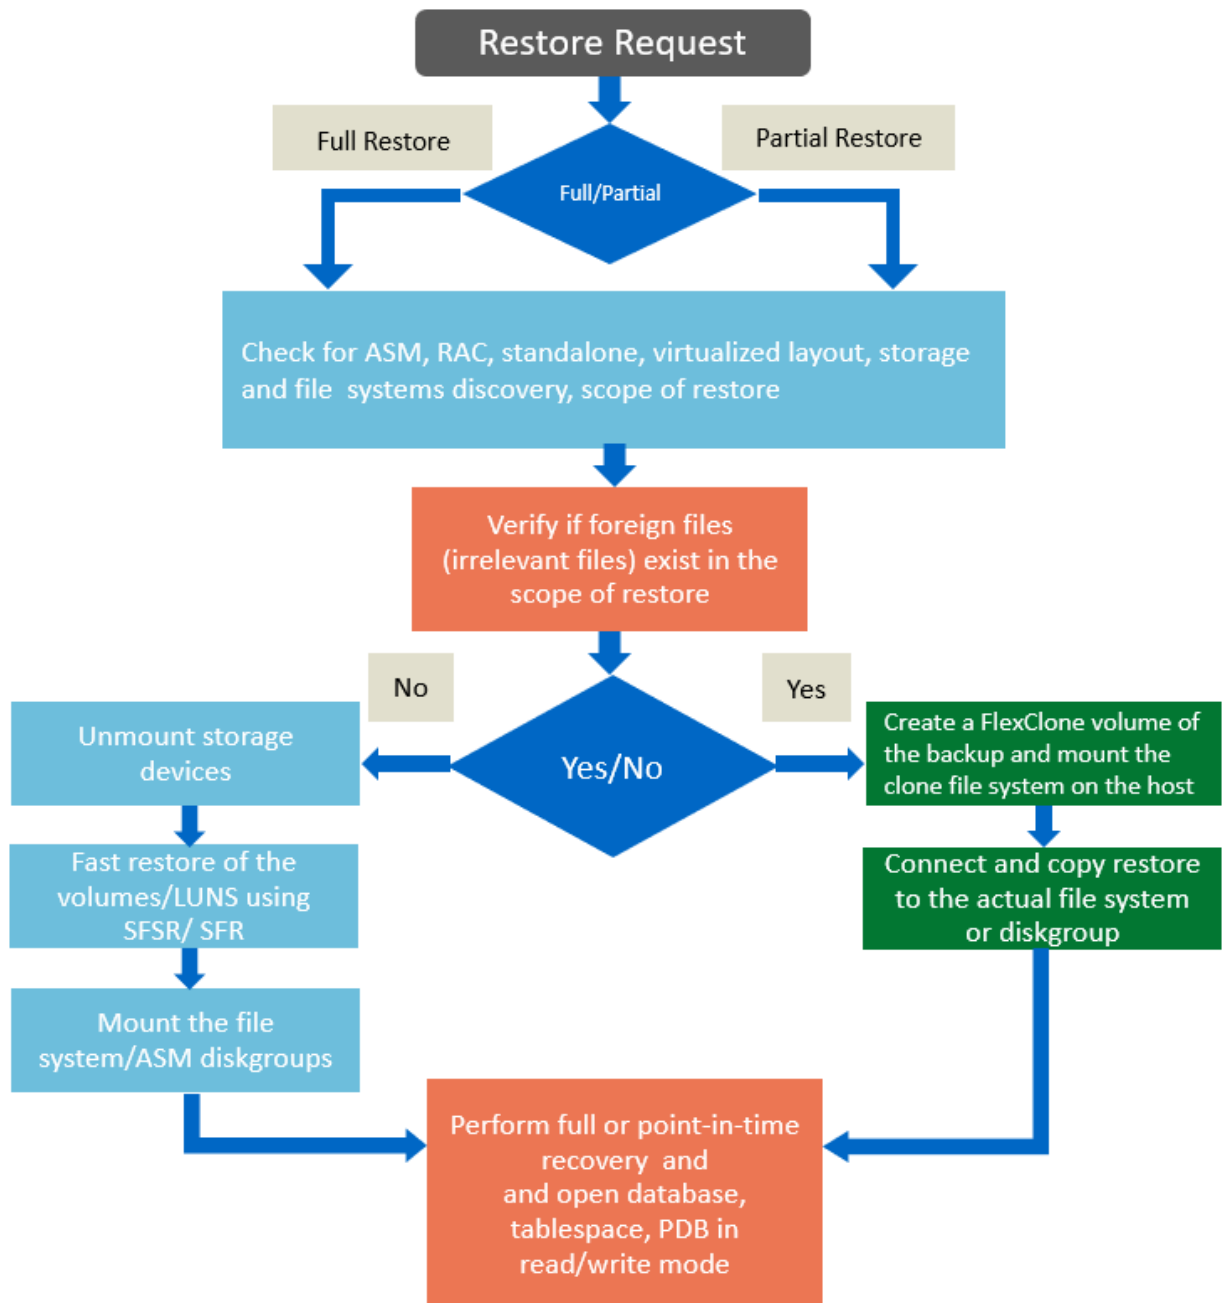
Additional Resources

For more information on how to backup an Oracle database across a Hybrid Cloud using SnapCenter see the following [YouTube demonstration video](#).

8 Restore and Recovery Best Practices

The SnapCenter Plug-In for Oracle supports the restore and recovery of an Oracle database. The restore and recovery can be done at full database level or at a granular entities level as low as PDB, PDB tablespaces, legacy 11g tablespace, and control file in a database. The flowchart in Figure 3 shows how a restore works in SnapCenter.

Figure 4) SnapCenter restore



The restore process is attempted either by using the storage Snapshot restore mechanism or through Oracle RMAN copy (connect and copy approach). The Snapshot restore process uses the ONTAP SFSR/single-file restore (SFR) technology, which is significantly faster than the Oracle RMAN connect and copy approach.

Here are the cases where restore uses the connect and copy approach rather than Snapshot restore:

- A restore from secondary in which the ONTAP version is earlier than 8.3 and SFR is not supported.

- Critical files such as spfile and passwd files are considered as nonoverridable files and if these are part of ASM DATA diskgroup.
 - If SPL is not installed or SPL is down on the remote node in the case of a RAC database.
 - If there are any structural changes to ASM disk group such as addition or removal of disks after a backup is taken.
 - If multiple tablespaces share the file system or disk group and the user selects only one or a subset of tablespaces.
 - If the LUN path and LUN serial number change for the ASM disk group after a backup is taken. Can be overridden by using the force-in-place option.
 - In the case of a file system on SAN (physical or RDM), the following scenarios use connect and copy for restore:
 - If there are nested mount points and the user selects a tablespace in a child file system.
 - If some other file system is mounted inside the file system requested for restore.
 - If there are multiple LVs and the user selects a tablespace residing on one of the LVs.
 - In conflicting restore modes such as:
 - Nested mount points where child is eligible for in-place restore but parent is doing connect and copy then connect and copy approach is chosen for both the cases
 - Multiple file systems in a single volume group where one file system uses connect and copy; but the other uses in-place restore, then connect and copy is chosen for all file systems in that volume group.
 - If LUN path and LUN serial number change for volume group after a backup is taken. It can be overridden by using the force-in-place option.
 - If there are any structural changes to the volume group such as addition or removal of LVs after a backup is taken.
- Note:** This can however be overridden by using the force-in-place option.
- In the case of NFS (VMDK over NFS datastore) datastore, the following scenario uses connect and copy for restore:
 - A restore from secondary in which the ONTAP version is earlier than 8.3 where SFR is not supported.
 - In the case of a file system on VMDK (VMFS datastore), the following scenario uses an approach similar to connect and copy restore:
 - Storage vMotion is used to copy VMDK from the cloned VMFS datastore to the actual datastore.

Once the restore activity is complete, you can recover the databases to a specific point in time, system change number (SCN), or the latest log present in the active file system. When recovering to an SCN or a time, SnapCenter checks for the archive logs present in the archive log or flash recovery area (FRA) destination and applies it for recovery.

8.1 Archive Log Management for Advanced Recovery

If the archive logs are not present in the active file system, then you must mount the log backup on the Oracle host and pass that location as an external log location in the restore recovery wizard. Failing to mount the required log backups might cause the entire recovery operation to fail. During the mount operation of a given backup, SnapCenter fires a FlexClone volume at the storage layer and mounts the FlexClone volume/LUN to the host.

To optimize the recovery time, pass multiple log backup destinations (the log backup that was mounted on the Oracle host) in the recovery wizard in ascending order of time, from oldest to newest. When the restore operation is complete, unmount the log backup.

To Mount the Log Backup and Pass the Mounted Destination

To mount the log backup and pass the mounted destination by using the recovery wizard, follow these steps.

1. Select the latest log backup and click the mount icon.

Manage Copies

15 Backups
0 Clones
Local copies

3 Backups
0 Clones
Mirror copies

10 Backups
0 Clones
Vault copies

Summary Card

28 Backups
11 Data Backups
17 Log Backups
0 Clones

Primary Backup(s)

Backup Name	Type	LF	End Date	Verified	Mounted	RMAN Cataloged	SCN
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-25-2018_18.44.02.1986_1	Log		4/25/2018 6:44:15 PM	Not Applicable	False	Cataloged	10694084
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-25-2018_17.44.03.2697_1	Log		4/25/2018 5:44:14 PM	Not Applicable	False	Cataloged	10691720

2. In the mount wizard, select the host and click the mount button.

If you are mounting from a secondary storage location, select the vault or mirror destination.

Mount backups

Choose the host to mount the backup

RHEL3.demo.netapp.com

Mount path : /var/opt/snapcenter/sco/backup_mount/RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-25-2018_18.44.02.1986_1/SCMPROD

Mount Cancel

3. Use the `df -h` command on the Oracle host to check for the mounted log backup.

```
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/vg_rhel1-lv_root              20G   7.1G   12G  38% /
tmpfs                                     3.9G  148K   3.9G   1% /dev/shm
/dev/sda1                                 485M   40M   421M   9% /boot
db_nfs_l1f1:/dr_oradata_pdb              5.7G   4.2G   1.6G  73% /oradata
db_nfs_l1f1:/dr_archive_pdb              12G   7.0G   4.8G  60% /archive
db_nfs_l1f1:/dr_rman_stage                8.8G   5.7G   3.2G  65% /oracle_home
db_nfs_l1f1:/Scs9f4e1a4-9269-458e-a33a-383db44e3631 5.7G   3.3G   2.5G  58% /var/opt/snapcenter/sco/backup_mount/RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_13.24.45.7137_0/SCMPROD/1
db_nfs_l1f1:/Sc14dbe2be-6adb-4afc-be97-8d6aee751972 12G   115M   12G   1% /var/opt/snapcenter/sco/backup_mount/RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_13.24.45.7137_1/SCMPROD/1
192.168.1.11:/Sc0533bb09-3dae-463e-b29c-e7488f58824b 6.9G   437M   6.5G   7% /var/opt/snapcenter/sco/backup_mount/RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-25-2018_08.12.07.3277_1/SCMPROD/1
192.168.1.11:/Sc0c7a36eb-5ab5-4601-85a4-d554f7862d17 6.9G   346M   6.6G   5% /var/opt/snapcenter/sco/backup_mount/RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-24-2018_08.12.07.1917_1/SCMPROD/1
192.168.1.11:/Sc173bbafd-65e2-4e7e-9de9-6f533464d461 6.9G   216M   6.7G   4% /var/opt/snapcenter/sco/backup_mount/RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-25-2018_17.44.03.2697_1/SCMPROD/1
db_nfs_l1f1:/Sc0ac6327f-20ee-4cad-b750-535028bf00ad 12G   254M   12G   3% /var/opt/snapcenter/sco/backup_mount/RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-25-2018_18.44.02.1986_1/SCMPROD/1
[oracle@rhel13] [SCMPROD] [-]#
```

4. Pass this mounted log location in the external archive log files location of the recovery wizard.

Restore SCMPROD

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

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

Previous

Next

- If you want multiple logs to be replayed for recovery, mount all of those log backups to the Oracle host and click the plus symbol to pass all these locations in the similar way.

Restore SCMPROD

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

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

6. Use the following formats to specify the external archive log location based on FRA and non-FRA layouts:
 - For a non-ASM file system, if the archive log destination is not configured on FRA, use the format `/mounted_archive_log/dbname/arch_dest`. Stop with the directory and do not specify the archive log file name.
 - For a non-ASM file system, if the archive log destination is configured on FRA, use the format `/mounted_archive_log/dbname/archivelog`. Notice the directories in the date format `YYYY_MM_DD` that contain the list of archive log files for the respective date in this directory.
 - For an ASM file system, if the archive log destination is configured on FRA, use the disk group format `+ <Mounted_Archivelog_diskgroup_name>`. Do not specify any other directories or files under this disk group name.
7. For non-FRA-based layout, you should pass the external archive log location in ascending order of the file or time; that is, in the order of the logs that the database needs for recovery. If not passed in the correct order, recovery might take a little longer than usual.
8. In the RAC restore use case, manually bring up the other RAC node after a successful restore on one of the nodes.

Note: You can do all the supported SnapCenter recovery operations in RMAN by cataloging the Snapshot copies with Oracle RMAN. You can catalog the metadata to target control file or catalog control file. The catalog control file settings can be edited in the database settings of the resource page.

8.2 Restore from Secondary Mirror or Vault Storage

As described in section 6, SnapCenter Policies and Resource Group Best Practices, there are two options to protect backups for near term retention in secondary storage, in a vault and/or disaster recovery destination. To understand the restore use cases from secondary, it is important to understand why the Snapshot copy must be replicated to secondary storage. Here are the few cases:

- Snapshot copies in primary storage are useful for shorter recovery time objectives (RTO). Mission-critical systems with narrow RTOs demand efficient RPOs, where backups are taken at very frequent intervals. However, this situation might cause the limit of 255 Snapshot copies to be reached, so retaining more backups on primary storage might not be possible. In such cases, you should replicate daily or weekly backups to secondary vault and/or disaster recovery storage.
- Replicating to a disaster recovery site (mirror destination), that is, mirroring all Snapshot copies supported by SnapCenter, is useful for bringing up a database in a secondary site in the event of a disaster.
- Vaulting a Snapshot copy helps to achieve a solution similar to tape. It is much faster than tape for restores and cloning, but not for storing multiple backups for more than 7 years (with optimum method of storing daily, weekly, monthly, and yearly). You can perform clones and restores with similar performance to that of the primary.

- Note:** For vault, mirror, and unified replication, you can perform restores and clones from both SnapCenter and storage. However, if protection is done outside of SnapCenter, where you have taken Snapshot copies directly from primary storage and enabled schedules to replicate them to the secondary destination, it might be necessary to manually perform all the steps for restore or clone from storage; in this case you cannot leverage SnapCenter.
- In SnapCenter, you can choose backups between primary and secondary storage for restores. Here are two reasons to choose backups from the mirror or vault destination for restore and recovery:
 - The availability of storage or storage failure at the production site.
 - The availability of Snapshot copies on the primary storage; that is, if Snapshot copies are already deleted on the primary storage based on the retention settings.

Here is a snippet of the topology view representing the backups located at different storage destinations (primary/local, secondary vault, and secondary mirror), with the vault destination selected.

The screenshot shows the 'SCMPROD Topology' interface. On the left, a 'Manage Copies' diagram shows a hierarchy: 'Local copies' (13 Backups, 0 Clones) branching into 'Mirror copies' (1 Backup, 0 Clones) and 'Vault copies' (10 Backups, 0 Clones). The 'Vault copies' are selected. On the right, a 'Summary Card' displays: 24 Backups, 11 Data Backups, 13 Log Backups, and 0 Clones. Below this is a table of 'Secondary Vault Backup(s)'.

Backup Name	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
RHEL3_demo_netapp.com_SCMPROD_RHEL3_04-25-2018_08.12.07.3277_1	Log	4/25/2018 8:12:41 AM	Not Applicable	False	Cataloged	10669794
RHEL3_demo_netapp.com_SCMPROD_RHEL3_04-25-2018_08.12.07.3277_0	Data	4/25/2018 8:12:24 AM	Unverified	False	Cataloged	10669790
RHEL3_demo_netapp.com_SCMPROD_RHEL3_04-24-2018_08.12.07.1917_1	Log	4/24/2018 8:12:37 AM	Not Applicable	False	Cataloged	10604957
RHEL3_demo_netapp.com_SCMPROD_RHEL3_04-24-2018_08.12.07.1917_0	Data	4/24/2018 8:12:23 AM	Unverified	False	Cataloged	10604953
RHEL3_demo_netapp.com_SCMPROD_RHEL3_04-23-2018_14.05.44.2010_1	Log	4/23/2018 2:06:11 PM	Not Applicable	False	Cataloged	10547036
RHEL3_demo_netapp.com_SCMPROD_RHEL3_04-23-2018_14.05.44.2010_0	Data	4/23/2018 2:05:58 PM	Verified	False	Cataloged	10547032

To Restore from a Secondary Vault Destination

1. Select a data backup and click the Restore button.

The screenshot shows the 'Restore SCMPROD' wizard. The '1 Restore Scope' step is active. The 'Secondary storage location' is 'Snap Vault / Snap Mirror'. The 'Source Volume' is 'db.demo.netapp.com:dr_oradata_pdb' and the 'Destination Volume' is 'db-dr:dr_oradata_pdb_vault'. Under 'Restore Scope', 'All Datafiles' is selected. Under 'Database State', 'Change database state if needed for restore and recovery' is selected. Under 'Restore Mode', 'Force in place restore' is selected. The wizard has 'Previous' and 'Next' buttons at the bottom.

2. In the restore wizard, notice the mapping of source volume and destination volume in the secondary storage location. Because you selected Snapshot copies from the vault destination, the wizard has selected the vault destination volume. Similarly, if the Snapshot copy was selected from the mirror destination, the wizard would show mirror destination volumes. If you have both, you can change the destination volume from vault to mirror.
3. The Restore Scope window contains options:
If you have selected all data files plus the control file, it's a full database restore of the legacy 11g database or the entire 12c CDB database, including PDB.

If you have selected a multitenant database, you have the option to choose PDB. You can choose multiple PDBs for full restore and recovery but not PITR- or SCN-based recovery.

For multitenant databases restore, you can also choose PDB tablespaces for full restore and recovery. You can select multiple tablespaces within a single PDB but not from multiple PDBs at a time. PITR- and SCN-based recovery are not supported.

Note: Selecting or deselecting control file for restores is crucial during the restore operation. If tablespaces or data files are added to or deleted from the database after the backup, and if control files are not selected for restore, restore and recovery fail due to inconsistency in the current control file and backup control file. The best practice is to take a backup whenever there is a change in the database architecture.

4. Select the **Change Database State** checkbox to bring the database offline during a complete database restore if the database is up and running.
5. Use the **Restore Mode** option to enforce an in-place restore (faster restore mechanism), even if it can't meet fast restore requirements due to foreign file constraints. If foreign files like non-Oracle files or backup copies or files from a different Oracle database exist on the same LUN as the actual database to be restored, the default connect and copy approach to fast restore is overridden, thus removing all files that were newly created and not part of the regular backup. On the contrary, if any of the Oracle files of the actual database exist or share the same volume/LUN, fast restore is not performed despite the checkbox being selected. Therefore the connect and copy approach is used in such cases. For more information about the mechanism and importance of restore-friendly layouts, refer to section 9, Oracle Clone Best Practices.
6. The **Recovery Scope** window contains four options:
 - All Logs. This option applies to all the logs present in the active file system for recovery. If you have passed any external archive log file locations, it applies the logs present in that mounted location and then applies the remaining logs until the latest available in the active file system.

Restore SCMPROD

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Choose Recovery Scope

☒ All Logs

☐ Until SCN (System Change Number)

☐ Date and Time

☐ No recovery

Specify external archive log files locations

Previous Next

- Until SCN. This option checks through the logs in the active file system that have this SCN. If the log isn't present in the active file system, you must mount the log backups and pass the location

in the external log location. This option works only for restore of the entire database; it is not displayed in the case of tablespace, PDB tablespace, and PDBs.

- Date and Time. This option is like SCN except that it uses a specific date and time. It is also disabled for restore of any granular entities (tablespace, PDB).
 - No recovery. SnapCenter performs only the restore operation of all data files, tablespace, or PDBs; it doesn't perform recovery. This option is useful for administrators to perform manual recovery by using Oracle SQLPLUS/RMAN.
7. (Optional) In the PreOps window, specify any scripts that you want to run before the restore operation.

The screenshot shows the 'Restore SCMPROD' window with the 'PreOps' tab selected. On the left is a vertical sidebar with six steps: 1 Restore Scope, 2 Recovery Scope, 3 PreOps (highlighted), 4 PostOps, 5 Notification, and 6 Summary. The main area is titled 'Specify optional scripts to run before performing a restore job' with an information icon. It contains three input fields: 'Prescript full path' with a text box containing '/var/opt/snapcenter/spl/scripts/' and a placeholder 'Enter Prescript path'; 'Arguments' with an empty text box; and 'Script timeout' with a numeric input set to '60' and a unit dropdown set to 'secs'.

8. (Optional) In the PostOps window, provision any scripts that should be run after the restore and recovery operation.
9. Select the checkbox to open the database or container database in read-write mode; otherwise you must open the database manually.
10. In the Notification window, configure email alerts for successful and failed operations.
11. Click Finish to submit the job.

Note: You can track the progress of the job in the monitor page or activity panel.

8.3 Disaster Recovery

While SnapCenter doesn't support an orchestrated disaster recovery solution, you can still use SnapCenter backups to perform manual disaster recovery. This section describes the steps for manual recovery of a database in the disaster recovery site.

1. If storage alone fails on site A:

You must break the SnapMirror relationship manually for each volume of the database, using the latest application-consistent Snapshot. You then discover those LUNs as ASM disks, mount them back as ASM disk groups on the host, and then bring up the database.

2. If an entire site fails such as storage, compute, host, or network:

You must break the SnapMirror relationship across the storage layer, using the latest application-consistent Snapshot copy, and bring up the storage on the disaster recovery site.

NetApp recommends keeping the disaster recovery Oracle host ready (in passive mode) to host disaster recovery volumes or LUNs directly. You could keep a similar compute machine or host with Oracle and grid home patched exactly the same as production. Verify that the network layer changes are handled effectively so that production traffic is redirected to the disaster recovery site. In the case of an ASM-

managed production database, NetApp recommends keeping the plain ASM instance up and running passively on the disaster recovery site.

For SAN-based deployment, set the `iscsiadm` connection or FC LUNs to be zoned, so that hosts can discover the LUNs from the storage. When these storage devices are discovered on the host (after the SnapMirror relationship has been broken using the latest application-consistent Snapshot), they must be mounted as ASM Diskgroup. Run the recovery command (PITR or SCN based) to bring up the database on the disaster recovery site. For recovery, check that all of the archive logs required for recovery are available on the active filesystem. If they are not, mount the log Snapshot copies to the host/ASM and catalog them with RMAN for automated recovery.

For NFS layouts, you can directly mount the NFS volumes and run the recovery command by replaying the logs and bring up the database.

If the disaster recovery Oracle host is not ready on the disaster recovery site, you must create a new host and have Oracle grid home (if it is ASM or RAC based) extracted from the source TAR backup. It might be necessary to relink the binaries with the new host. Once the Oracle grid home is ready, repeat step 2 to bring up the database by mounting the storage devices to the host or ASM, recreate the disk group, mount them to the ASM instance, and recover the database.

Another alternative is to bring up a clone of the Oracle production database from a secondary disaster recovery or vault-replicated Snapshot. To host a clone to the disaster recovery host, first make sure that the SnapCenter agent for Oracle was installed on that host. Second, Oracle home must have been configured. Third, if the source was an ASM database, the grid home and ASM instance must be up and running on that host. The clone SID can still use the same name as the production database. Once the clone is completed, you can split the clone from the vaulted volumes to get standalone volumes for the database.

8.4 Important Considerations

If the connect and copy restore method takes too long to finish, large restores might fail. In such cases, the workaround is to edit `/var/opt/snapcenter/sco/etc/sco.properties` and modify the following line:

```
ORACLE_PLUGIN_RMAN_TIMEOUT=72
```

After doing this, restart the SPL service with `/opt/NetApp/snapcenter/spl/bin/spl restart`.

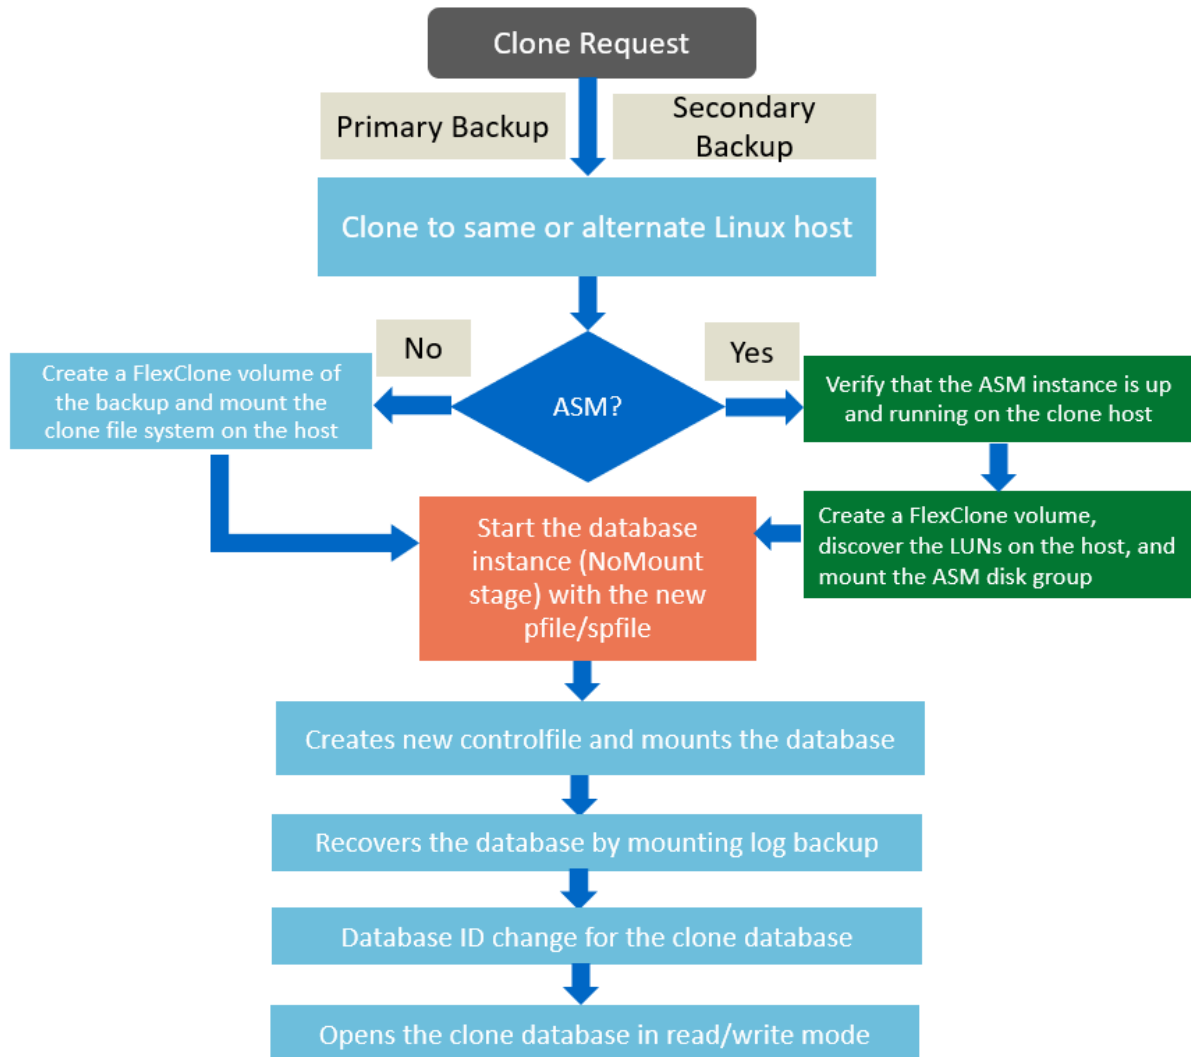
Additional Resources

For more information on how to perform a quick and easy restore of Oracle multitenant database using NetApp SnapCenter, see the following [YouTube demonstration video](#).

9 Oracle Clone Best Practices

Figure 1 is a workflow for a SnapCenter request to clone an Oracle database.

Figure 5) SnapCenter clone



You can use SnapCenter to clone a database by using the backup of the database. The clone operation creates a copy of the database data files, and creates new online redo log files, control files, and archive log destination. The clone database can also be recovered to a specified time, SCN, or Until Cancel.

Here are the prerequisites or best practices before you perform a clone operation:

- Create a backup of the database by using SnapCenter.
Create either an online data and log backups or an offline (mount or shutdown) backup.
- If you want to customize the control file or redo log file paths, verify that you have reprovisioned the required file system or ASM disk group.
By default, redo log and control files of the cloned database are created on the ASM disk group or the file system provisioned by SnapCenter for the data files of the clone database.

- A clone can be created on the same host as that of the source database. When creating the clone on an alternate host, confirm that the alternate host meets the following requirements:
 - SnapCenter Plug-In for Oracle database is installed on the alternate host.
 - The clone host is able to discover LUNs from primary or secondary storage.
 - If you are cloning from primary storage or secondary (vault or mirror) storage to an alternate host, make sure that an iSCSI session is established between the primary or secondary storage and the alternate host, or that it is zoned properly for FC.
 - If you are cloning from vault or mirror storage to the same host, make sure that an iSCSI session or FC zoning is established between the vault or mirror storage and the Oracle host.
 - If you are cloning in a virtualized environment, verify that an iSCSI session is established between the primary or secondary storage and the ESX server hosting the alternate host, or that it is zoned properly for FC.
 - The Oracle version is the same as that of the source database host.
 - The OS distribution and version are the same as those of the source database host.
 - If you want to override the clone to a different operating system version, specify the parameter as follows:
 Set the parameter `ALLOW_CLONE_OS_MISMATCH=TRUE` in the `/var/opt/snapcenter/sco/etc/sco.properties` file.
 Restart the plug-in service `/opt/NetApp/snapcenter/spl/bin/spl restart`.
- If the source database is an ASM database:
 - Make sure that the ASM instance is up and running on the host where the clone is being performed.
 - If you want to place archive log files of the cloned database in a dedicated ASM disk group, make sure that the ASM disk group is provisioned prior to the clone operation.
 - The name of the data disk group can be configured, but you should verify that the name is not used by any other ASM disk group on the host where the clone is being performed. Data files residing on the ASM disk group are provisioned as part of SnapCenter clone workflow.
- To clone a backup of a 12c database, set the value of `exclude_seed_cdb_view` to `FALSE` in the source database parameter file to retrieve seed PDB-related information.
 The seed PDB is a system-supplied template that the CDB can use to create PDBs. The seed PDB is named `PDB$SEED`. For information about `PDB$SEED`, see the Oracle Doc ID 1940806.1.
- You can also perform a clone operation from Data Guard and Active Data Guard Standby databases. Doing so avoids using the production volumes or database host for cloning. Remember to follow backup best practices for Data Guard and Active Data Guard configurations outlined in section 7, SnapCenter Resource Group and Policies Best Practices.

Cloning is a straightforward approach in which a specific data Snapshot is selected from either primary or secondary replicated storage.

9.1 To Perform a Clone Operation by Using the Clone Wizard

1. In the clone wizard, enter the clone SID.

By default, the clone wizard populates the storage volume mapping for the source volumes of the given production database. If you have more than one mapping — that is, both SnapMirror and SnapVault — you can still change the destination volume to either vault or mirror volume.

Clone from SCMPROD

1 Name

Provide clone database SID

Clone SID

Secondary storage location : Snap Vault / Snap Mirror

Data

Source Volume	Destination Volume
db.demo.netapp.com:dr_oradata_pdb	db-dr:dr_oradata_pdb_vault

Logs

Source Volume	Destination Volume
db.demo.netapp.com:dr_archive_pdb	db-dr:dr_archive_pdb_mirror

Previous Next

2. In the Locations window, enter the clone host details.

The clone host can be the same host as that of production, or it can be a different Linux host. If you plan to clone to the same or an alternate host, you must follow the prerequisites. You can customize the directory structure and default values populated for data files, redo log, and control file. If you are cloning in a hybrid cloud environment, the clone Linux host can be in the cloud, like AWS, Azure, IBM, and so on. The network firewall must be open to listen to the host IPs and the ports. It lives in the primary and secondary storage SVMs that were listed in the preinstallation guidelines in section 4, Preinstallation Best Practices for SnapCenter.

Clone from SCMPROD

1 Name

Provide clone database SID

Clone SID

Secondary storage location : Snap Vault / Snap Mirror

Data

Source Volume	Destination Volume
db.demo.netapp.com:dr_oradata_pdb	db-dr:dr_oradata_pdb_vault

Logs

Source Volume	Destination Volume
db.demo.netapp.com:dr_archive_pdb	db-dr:dr_archive_pdb_mirror

Previous Next

3. In the Credentials window, enter the details for hosting the clone database: Oracle home, OS user, and group of the target host. Enter the system credentials for the clone database.

Clone from SCMPROD

1 Name
2 Locations
3 Credentials
4 PreOps
5 PostOps
6 Notification
7 Summary

Database Credentials for the clone

Run As name for sys user + ⓘ

Oracle Home Settings ⓘ

Oracle Home

Oracle OS User

Oracle OS Group

- (Optional) In the PreOps window, you can provide prescripts that can be run before executing the clone. Any executable scripts, such as Shell, Perl, and Python scripts, are acceptable, but they must be kept in the default location.
- (Optional) You can also enter database parameter settings, which are your pfile/spfile settings for the clone database. By default, all the values for each parameter present in the production pfile/spfile are copied. You should add or edit the archive log destination for your clone database, or else the files are placed in the ORACLE_HOME destination. Similarly, you should check the SGA, PGA, and open cursors values for the clone database.
- (Optional) In the PostOps window, you can opt out of recovery. If you have selected any of the recovery options, SnapCenter automates mounting the immediate archive log backup after the data backup on the target clone host.

Clone from SCMPROD

1 Name
2 Locations
3 Credentials
4 PreOps
5 PostOps
6 Notification
7 Summary

☒ Recover Database

☒ Until Cancel ⓘ

☐ Date and Time ⓘ

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

☐ Until SCN (System Change Number) ⓘ

Specify external archive log locations ⓘ ⓘ

Until Cancel. This option applies all the logs in the mounted log backups and brings up the database. By default, it applies logs only from the log backup that was automatically mounted by SnapCenter. For example, to recover to the latest log backups that were taken, mount the log backups manually to the clone host and pass those locations in the external archive log locations.

Until Date and Time or SCN. If you pass date and time or SCN, check that the log backup that is mounted by SnapCenter holds good for that recovery point. If it's part of a different log backup, mount the respective backups on the clone host and pass those locations in external archive log locations.

Note: You can also pass live SQL queries in the PostOps window, along with regular scripts. For example, to delete all HR-sensitive bank or password records, you can write a wrapper shell

or Perl script and pass it as postscript for the clone. You can also use the clone wizard or the CLI to automate an end-to-end clone.

9.2 To Perform a Clone Operation by Using the CLI

1. Create an Oracle database clone specification from a specified backup.

The command automatically creates an Oracle database clone specification file for the specified source database and its backup. You must also enter a clone database SID so that the specification file created has the automatically generated values for the clone database that you are creating. You can also specify the recovery options, the host where the clone operation is to be performed, prescripts, postscripts, and other details.

```
sccli New-SmOracleCloneSpecification -AppObjectId [-BackupName | -CloneLastBackup ] -CloneDatabaseSID [-IncludeSecondaryDetails] [-SecondaryStorageType ] [- SetConsoleOutputWidth]
```

Sample Syntax

```
[root@rhel-linux ~]# sccli New-SmOracleCloneSpecification -AppObjectId 'rhellinux.netapp.com\STddb' -CloneLastBackup 2 -CloneDatabaseSID 'CDBCLONE'
```

INFO: You have chosen to generate clone specification using last backup number '2' having backup name 'federated-ds_rhel-linux_10-25-2015_22.30.30.4523_0'.

INFO: Oracle clone specification file
'/var/opt/SnapCenter/sco/clone_specs/oracle_clonespec_CDB_CDBCLONE_2015-10- 25_23.59.12.317.xml'
got created successfully.

INFO: The command 'New-SmOracleCloneSpecification' executed successfully..

2. Initiate a clone operation from an existing backup.

This command initiates a clone operation. You must also enter an Oracle clone specification file path for the clone operation.

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

Sample syntax

```
[root@rhel-linux ~]# sccli New-SmClone -CloneToHost 'rhel-linux.netapp.com' -OracleCloneSpecificationFile '/var/opt/snapcenter/sco/clone_specs/oracle_clonespec_CDB_CLONE12C_2015-11-26_00.20.29.237.xml'
```

INFO: Recovery of the cloned Oracle Database will be performed using all available logs in immediate log backup after the data backup chosen for clone because neither SCN nor time is specified.

INFO: Job 'Clone from backup 'stddb-ds_rhel-linux_11-24-2015_00.55.10.2377_0'' QUEUED with jobId '364'

INFO: The command 'New-SmClone' executed successfully.

For more information about the CLI commands, refer to the [Command Reference Guide](#).

9.3 RAC-to-RAC Clone

A RAC database is cloned as a standalone database by using SnapCenter. It is also possible to convert a non-RAC clone database to a RAC database. It is assumed that the second RAC node is already part of the cluster.

To Convert a Non-RAC Clone Database to a RAC Database

1. Use the GUI or CLI (non-RAC clone) to perform a regular clone operation.
2. Create redo and undo for the second instance:

```
alter database add logfile thread 2 group 3 ('+DATA','+FLASH') size 50m reuse;
```



```
alter database add logfile thread 2 group 4 ('+DATA','+FLASH') size 50m reuse;
alter database enable public thread 2;
create undo tablespace UNDOTBS2 datafile '+DATA' size 50G;
```

3. Add cluster-related parameters in the `init<sid>.ora` file:

```
*.cluster_database_instances=2
*.cluster_database=true
*.remote_listener='LISTENERS_ORCLDB'
ORCLDB1.instance_number=1
ORCLDB2.instance_number=2
ORCLDB1.thread=1
ORCLDB2.thread=2
ORCLDB1.undo_tablespace='UNDOTBS1'
ORCLDB2.undo_tablespace='UNDOTBS2'
#update the actual controlfile path
*.control_files='+DATA/ORCLDB/controlfile/current.256.666342941','+FLASH/ORCLDB/controlfile/curre
nt.256.662312941'
```

4. Copy the updated `init.ora` file to node2 and rename the files as per the instance name:

```
[oracle@orac1]$ mv initORCLDB.ora initORCLDB1.ora [oracle@orac2]$ mv initORCLDB.ora
initORCLDB2.ora
```

5. Register the cloned RAC database with `srvctl`:

```
[oracle@orac1]$ srvctl add database -d ORCLDB -o /u01/app/oracle/product/12.2/db_1
[oracle@orac1]$ srvctl add instance -d ORCLDB -i ORCLDB1 -n orac1
[oracle@orac1]$ srvctl add instance -d ORCLDB -i ORCLDB2 -n orac2
```

6. Stop and start the services by using `srvctl` and perform a sanity check by using `crsctl` command
7. All the steps performed to turn this clone into a RAC database must be undone before attempting to delete the clone.

9.4 Important Considerations

- In the case of a load-sharing (LS) mirror, the clone operation might fail if LS mirror updates take too long on the storage system. The workaround is to include the following parameters and specify the value in the `<appSettings>` section of the `SMCoreServiceHost.exe` configuration file, located under `SMCore` in the SnapCenter Server. Then restart the SnapCenter `SMCore` service.

```
<add key="lsmsleep" value="300000">
```

The value 300000 (in ms) results in a 300-second wait.

- When a clone from SnapVault secondary for FC SAN/ASM configuration fails with `Error` executing SQL `"ALTER DATABASE OPEN RESETLOGS` within 2100 seconds against Oracle database, the workaround is to run the following commands on the Linux node that is hosting the clone:

```
cd /opt/NetApp/snapcenter/spl/bin
./sccli Open-SmConnection
./sccli Set-SmConfigSettings -ConfigSettingsType Plugin -PluginCode SCO -ConfigSettings
"KEY=ORACLE_SQL_QUERY_TIMEOUT,VALUE=10800"
./sccli Set-SmConfigSettings -ConfigSettingsType Plugin -PluginCode SCO -ConfigSettings
"KEY=ORACLE_PLUGIN_SQL_QUERY_TIMEOUT,VALUE=10800"
```

- Restart the SPL process to reflect the change: `./spl restart`

Additional Resources

For more information about using SnapCenter to clone an Oracle database to an alternate host running in a public cloud (AWS, Azure), see the following [YouTube demonstration video](#).

Appendix

A. SnapCenter Deployment Models for Oracle Database

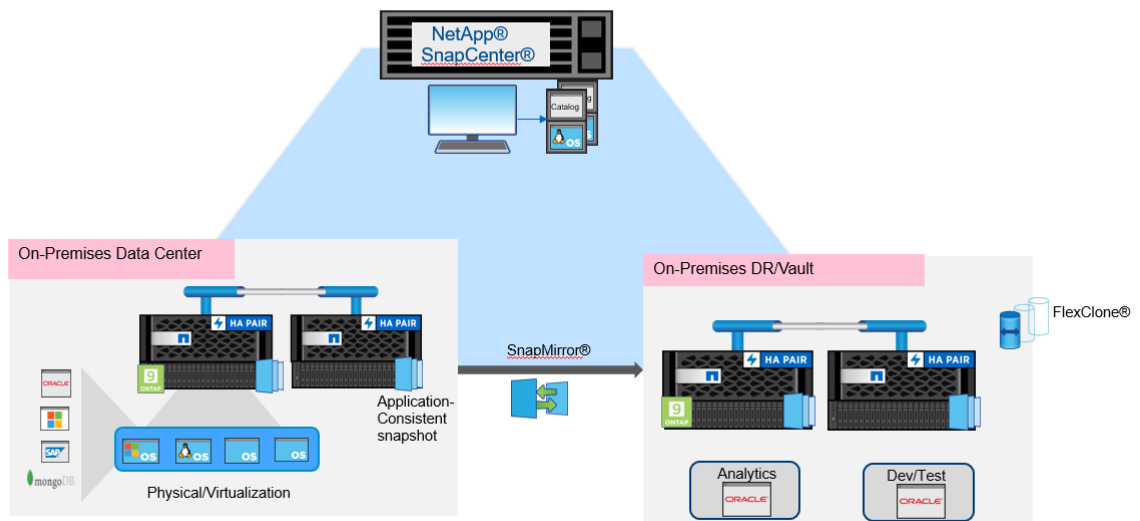
SnapCenter can be deployed in either of three models:

- Private data center/cloud
- Public cloud
- Hybrid cloud

Private data center/cloud. In this deployment, both the primary and secondary vault or disaster recovery storage Oracle hosts are running on premises. This is the most generic model.

Figure 6) Private data center/cloud deployment

On-premises Disaster Recovery/Vault deployment



Hybrid cloud deployment using a collocated data center is also called next-to-cloud deployment using NetApp Private Storage. In this deployment, the secondary storage is completely running in a private (controlled) environment hosted in a collocated data center, such as Equinix, and the compute is leveraged from the cloud. This is probably the safest model, because the data is in your control. The advantage of this model is that you can switch between different hyperscalar clouds, based on the workload requirements against cost.

Figure 7) NetApp Private Storage deployment

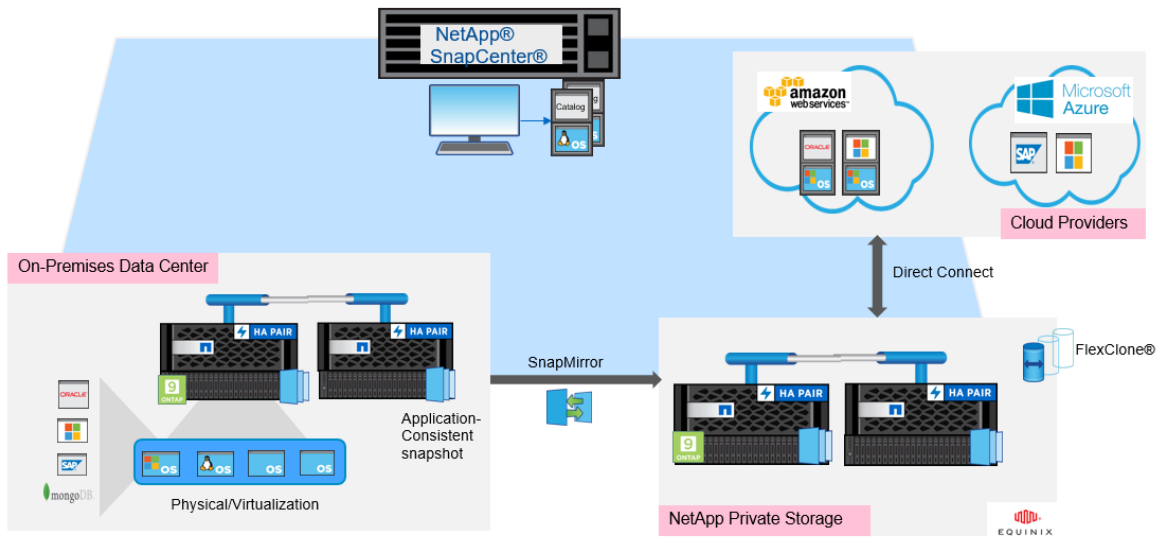
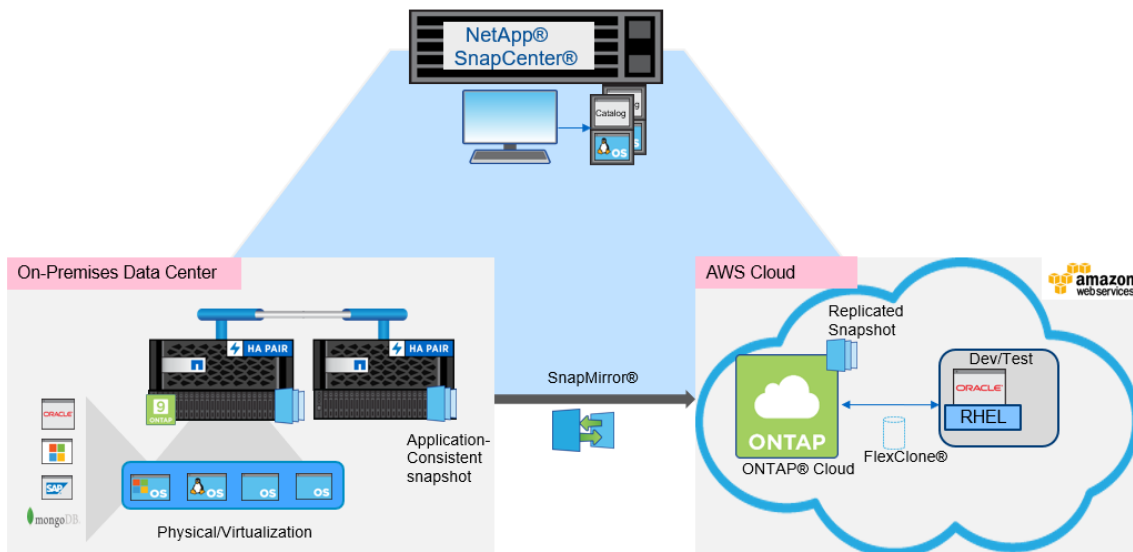


Figure 7 shows a typical hybrid cloud deployment. The primary storage is running on premises and secondary storage (vault or disaster recovery) is running on a public cloud (Amazon, Azure, and so on). The ONTAP cloud storage is running on public cloud as a replica destination. This is very a good reference model for dev/test scenarios. You can dynamically spin hosts and host a dev/test whenever required. This model also avoids holding any dedicated infrastructure for secondary disaster recovery or vault in the on-premises data center.

Figure 8) Hybrid cloud deployment



B. RAC One Node and Other Third-Party Cluster Solutions (Active-Passive)

RAC One Node and other active-passive cluster solutions works well with SnapCenter. During failover of the database to an alternate cluster node it is required to follow the below guidelines to resume the operations without any interruption..

To Work with RAC One Node Layouts

1. Register all the nodes in the Oracle RAC one Node cluster.

For example : In a 2-node scenario where Node 1 is active and Node 2 is passive. Register all the nodes with SnapCenter like an Oracle RAC hosts.

2. Discover the databases on all the nodes.

Note: Make sure the oracle database entry is added in the /etc/oratab of the Node2(Passive) to support Database failover scenarios.

The Rac One Node database is discovered as the RAC database on the node (Node1) where it is currently hosted. The same database is discovered as a standalone database on all the other RAC nodes (In this Example:- Node2). The Database name listed from Node2 can be considered as stale as long as failover happened to second node.

3. All the workflows(Backup, restore and clone) tested on Node1 should be successful.
4. When you have migrated the Rac One Node database from Node1 to Node2 then perform a manual resource discovery on both nodes by refreshing the resource page.

The RacOneNode database is discovered as a RAC database on the node (Node2) where it is currently hosted. The same database is discovered as a standalone database on all the other RAC nodes (In our case, Node1).

All the backups performed previously on other nodes of RAC One Node are still visible and are valid.

5. When the database is migrated to Node 2 from Node 1, you must manually configure the preferred host as Node2 in preferred RAC Node (Configure database settings)
6. You can exercise restore and clone workflows in the usual way on the active node.
7. Similarly, you can migrate back the Rac One Node database from Node2 to Node1 and perform a manual resource discovery on both nodes. As mentioned in step 4, you must manually configure the preferred host as Node1.

Active-Passive Cluster Solutions

For other cluster solutions (active–passive) like Red Hat Cluster Suite and SIOS follow the below guidelines. For example, if Oracle database is running on a Red Hat cluster that has cluster services running on two nodes, with one being active where the database is running and the other being passive.)

1. Install the plug-in manually on the Oracle hosts as root or non-root user with sudo privilege. Confirm that the SnapCenter Server IP and the SPL_ENABLED_PLUGINS have the correct values for the file: /var/opt/SnapCenter/spl/etc/spl.properties

```
SNAPCENTER_SERVER_HOST=10.232.206.110
SPL_ENABLED_PLUGINS=SCO,SCU
```

2. Use `sccli` to set the preferred IPs of all the cluster nodes and the host, including the floating IP; navigate to the below location to run `sccli` :-

```
root> cd /opt/NetApp/snapcenter/spl/bin
Sccli Set-PreferredHostIPsInStorageExportPolicy -IPAddresses
'10.231.73.50','10.231.73.51','10.231.73.52'
```

Here,

'10.231.73.51' is the floating IP that is being accessed by all the cluster nodes.

'10.231.73.52' is the public IP of the second cluster node, which is passive, and

'10.231.73.50' is the public IP address of the first node, which is actively hosting the database.

3. Repeat steps 1 and 2 on the second cluster node.
4. Add a local entry of floating IP '10.231.73.51' with a random hostname (FQDN) — for example, 'Linux-cluster.netapp.com' — in the Hosts file (C:\Windows\System32\drivers\etc) of the SnapCenter Server running on Windows. No node entries are required because the SnapCenter operations are completely carried over floating IP or VIP and therefore are not dependent on the host.
5. Use PowerShell or the GUI to perform the Add Linux Host operation and pass the 10.231.73.51 (floating IP) in the place of Host ip address. Select the checkbox to skip the preinstallation check and then complete the registration process. It might take some time to recognize the plug-in.
6. Refresh the resources screen to discover all of the Oracle databases running on Node 1 (where the current floating IP is running with databases; that is, the active node). You can also create policies and resource groups, start protecting them, and perform backups, restore, and clones.

During host failover, the Oracle database is migrated to an alternative node in the cluster that is using the floating IP or virtual IP. The backups and clones that were taken earlier on Node 1 are still reflected in SnapCenter, regardless of the database being migrated to the second node; it doesn't really matter because all of the workflows are completely communicated using only VIP.

Similarly, the databases can fall back to Node 1 and all backup, restore, mount, and clone functions still work normally.

C. Block-Level Recovery

To Perform Block-Level Recovery

You can perform block-level recovery by using these RMAN cataloged Snapshot copies. The following steps are an example of how to perform block level recovery.

1. Use the dbv utility to check the data file that has a corrupt block.

```
[oracle@orcldev114 bin]$ ./dbv file=/DATA2/PAYB/appsbiz01.dbf
DBVERIFY: Release 12.1.0.2.0 - Production on Thu Nov 2 12:01:06 2017
Copyright (c) 1982, 2014, Oracle and/or its affiliates. All rights reserved.
DBVERIFY - Verification starting : FILE = /DATA2/PAYB/appsbiz01.dbf
Page 8 is marked corrupt
Corrupt block relative dba: 0x01400008 (file 5, block 8)
Bad check value found during dbv:
Data in bad block:
  type: 30 format: 2 rdba: 0x01400008
  last change scn: 0x0000.007fb941 seq: 0x1 flg: 0x04
  spare1: 0x0 spare2: 0x0 spare3: 0x0
  consistency value in tail: 0xb9411e01
  check value in block header: 0x818a
  computed block checksum: 0x3b31
DBVERIFY - Verification complete

Total Pages Examined      : 6400
Total Pages Processed (Data) : 4
Total Pages Failing (Data) : 0
Total Pages Processed (Index): 0
Total Pages Failing (Index): 0
Total Pages Processed (Other): 129
Total Pages Processed (Seg) : 0
Total Pages Failing (Seg) : 0
Total Pages Empty         : 6266
Total Pages Marked Corrupt : 1
Total Pages Influx        : 0
Total Pages Encrypted     : 0
Highest block SCN         : 8999971 (0.8999971)
[oracle@orcldev114 bin]$
```

2. Use the following command to list the backups that are taken through SnapCenter and cataloged with RMAN:

```
RMAN> LIST DATAFILECOPY ALL;
```

```
List of Datafile Copies
```

```
=====
```

Key	File S	Completion Time	Ckp SCN	Ckp Time
341	1 A	08-NOV-17	9379266	08-NOV-17
Name: /var/opt/snapcenter/sco/backup_mount/orcldev114_nb_openenglab_netapp_com_PAYB_orcldev114_11-08-2017_02.28.03.3606_0/PAYB/1/PAYB/system01.dbf				
Tag: SCO_PAYB_1101				
331	1 A	08-NOV-17	9365506	08-NOV-17
Name: /var/opt/snapcenter/sco/backup_mount/orcldev114_nb_openenglab_netapp_com_PAYB_orcldev114_11-07-2017_20.28.03.7086_0/PAYB/1/PAYB/system01.dbf				
Tag: SCO_PAYB_1094				
323	1 A	07-NOV-17	9343866	07-NOV-17
Name: /var/opt/snapcenter/sco/backup_mount/orcldev114_nb_openenglab_netapp_com_PAYB_orcldev114_11-07-2017_14.28.02.9932_0/PAYB/1/PAYB/system01.dbf				
Tag: SCO_PAYB_1087				

3. Select a data file copy with an SCO prefixed tag that matches to the earliest recovery point and mount that data backup by using SnapCenter to the Oracle host for block-level restore.

Note: If logs on the active file system are already pruned, you can mount the respective log backups for recovery. Refer to section 8.1, Archive Log Management for Advanced Recovery, for information about how to use SnapCenter to mount the data and log backup.

4. Once the backups are mounted on the host, you can verify by running `df -h`. It should look similar to this:

```
10.195.48.151:/Sc8699dfac-6e46-43ba-9376-48d6b7bd7893
21G 2.2G 18G 11% /var/opt/snapcenter/sco/backup_mount/orcldev114_nb_openenglab_netapp_com_PAYB_orcldev114_11-01-2017_21.29.10.2574_0
10.195.48.151:/Scb6cc950a-ce4c-4cba-a3f1-ee71a46dc72b
21G 2.2G 18G 11% /var/opt/snapcenter/sco/backup_mount/orcldev114_nb_openenglab_netapp_com_PAYB_orcldev114_11-01-2017_21.29.10.2574_1
```

5. If the logs are not cataloged with RMAN, you have two options:

Copy the archived logs to the required destination and pass the location in the RMAN recovery.

Manually RMAN catalog the location or files of the mounted archive log destination.

When you have done this, invoke the recovery command:

```
RMAN> blockrecover datafile '/DATA2/PAYB/appsbiz01.dbf' block 8 from tag SCO_PAYB_1101;
```

```
Starting recover at 08-NOV-17
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=22 device type=DISK
```

```
channel ORA_DISK_1: restoring block(s) from datafile copy /var/opt/snapcenter/sco/backup_mount/orcldev114_nb_openenglab_netapp_com_PAYB
```

```
starting media recovery
```

```
archived log for thread 1 with sequence 243 is already on disk as file /DATA2/PAYB/archivelog/2017_11_02/o1_mf_1_243_dzody4om_.arc
archived log for thread 1 with sequence 244 is already on disk as file /DATA2/PAYB/archivelog/2017_11_02/o1_mf_1_244_dzooxr1_.arc
archived log for thread 1 with sequence 245 is already on disk as file /DATA2/PAYB/archivelog/2017_11_02/o1_mf_1_245_dzpc0w1x_.arc
archived log for thread 1 with sequence 246 is already on disk as file /DATA2/PAYB/archivelog/2017_11_02/o1_mf_1_246_dzpl1wyb_.arc
archived log for thread 1 with sequence 247 is already on disk as file /DATA2/PAYB/archivelog/2017_11_03/o1_mf_1_247_dzq045b5_.arc
archived log for thread 1 with sequence 248 is already on disk as file /DATA2/PAYB/archivelog/2017_11_03/o1_mf_1_248_dzq06vss_.arc
archived log for thread 1 with sequence 249 is already on disk as file /DATA2/PAYB/archivelog/2017_11_03/o1_mf_1_249_dzrb9t2j_.arc
archived log for thread 1 with sequence 250 is already on disk as file /DATA2/PAYB/archivelog/2017_11_03/o1_mf_1_250_dzrzd1lq_.arc
archived log for thread 1 with sequence 251 is already on disk as file /DATA2/PAYB/archivelog/2017_11_03/o1_mf_1_251_dzs6frd9_.arc
archived log for thread 1 with sequence 252 is already on disk as file /DATA2/PAYB/archivelog/2017_11_04/o1_mf_1_252_dzsnhwk2_.arc
archived log for thread 1 with sequence 253 is already on disk as file /DATA2/PAYB/archivelog/2017_11_04/o1_mf_1_253_dzt2jzjy_.arc
archived log for thread 1 with sequence 254 is already on disk as file /DATA2/PAYB/archivelog/2017_11_04/o1_mf_1_254_dzt9ncjz_.arc
archived log for thread 1 with sequence 255 is already on disk as file /DATA2/PAYB/archivelog/2017_11_04/o1_mf_1_255_dztyox6f_.arc
archived log for thread 1 with sequence 256 is already on disk as file /DATA2/PAYB/archivelog/2017_11_04/o1_mf_1_256_dzvfbb0t_.arc
```

- When the block recovery command is executed, you can use the dbv utility to check for corrupt blocks:

```
[oracle@orcldev114 bin]$ ./dbv file=/DATA2/PAYB/appsbiz01.dbf

DBVERIFY: Release 12.1.0.2.0 - Production on Wed Nov 8 12:53:29 2017

Copyright (c) 1982, 2014, Oracle and/or its affiliates. All rights reserved.

DBVERIFY - Verification starting : FILE = /DATA2/PAYB/appsbiz01.dbf


DBVERIFY - Verification complete

Total Pages Examined          : 6400
Total Pages Processed (Data)  : 4
Total Pages Failing (Data)    : 0
Total Pages Processed (Index): 0
Total Pages Failing (Index)   : 0
Total Pages Processed (Other): 130
Total Pages Processed (Seg)   : 0
Total Pages Failing (Seg)     : 0
Total Pages Empty             : 6266
Total Pages Marked Corrupt    : 0
Total Pages Influx             : 0
Total Pages Encrypted         : 0
Highest block SCN              : 8999971 (0.8999971)
```

D. Tablespace Point-in-Time Recovery

To Perform Tablespace Point-in-Time Recovery

The steps for tablespace point-in-time recovery are similar to those for block-level recovery; the only difference is that you must bring the tablespace offline before running the RMAN recovery command.

- Use the following RMAN script for tablespace point-in-time recovery.

```
RMAN> recover tablespace i2e until time "to_date('2012-06-07 12.03.00', 'YYYY-MM-DD HH24:MI:SS')"
auxiliary destination '/tmp'.

*
. importing SYS's objects into SYS
. . importing table "I2ET1"
Import terminated successfully without warnings.
host command complete
sql statement: alter tablespace I2E online
starting full resync of recovery catalog
full resync complete
sql statement: alter tablespace I2E offline
starting full resync of recovery catalog
full resync complete
sql statement: begin dbms_backup_restore.AutoBackupFlag(TRUE); end;
starting full resync of recovery catalog
full resync complete
Removing automatic instance
Automatic instance removed
auxiliary instance file /tmp/TSPITR_HRA_SFFT/onlinelog/ol_mf_3_7x0wvco6_.log deleted
auxiliary instance file /tmp/TSPITR_HRA_SFFT/onlinelog/ol_mf_2_7x0wvbjq_.log deleted
auxiliary instance file /tmp/TSPITR_HRA_SFFT/onlinelog/ol_mf_1_7x0wv93y_.log deleted
auxiliary instance file /tmp/TSPITR_HRA_SFFT/datafile/ol_mf_temp_7x0wvf9n_.tmp deleted
```

- After recovering the tablespace, bring it back on line and perform sanity checks.

E. Recover a Table

This section describes how to recover a corrupted, deleted, or dropped table.

To Recreate a Dropped Table by Exporting It from a Clone of a Backup

In this scenario, a table is dropped and must be imported back from an existing online backup. To restore only that table, first create a clone from the Snapshot backup by using SnapCenter. Then manually export the table from the clone database and manually import it back into the target database.

1. Use the GUI or CLI to create a clone of the target Oracle database on the same or a remote host (Refer to the steps in section 9, Oracle Clone Best Practices.)
2. When the clone is complete, manually export the table from the clone:

```
[oracle@tardb_host1][expl][~]$ exp userid=user/password tables=sales file=sales12.dmp
```

3. When the export is complete, manually import the table into the target database:

```
[oracle@tardb_host1][tardb1][~]$ imp userid=user/password tables=sales file=sales12.dmp
```

To Recreate a Dropped Table from a Clone of a Backup by Using a Database Link

In this scenario, a table is dropped and must be recreated from an existing online backup. To recreate just that table, first use SnapCenter to create a clone from the backup. Then manually create a database link from the target database to the clone and use the link to recreate the table in the target database.

1. Use the GUI or CLI to create a clone of the target Oracle database on the same or a remote host. (Refer to the steps in section 9, Oracle Clone Best Practices.)
2. When the clone is complete, manually add an entry for the clone database (for example, apr12cln) in the `tnsnames.ora` file.
3. Create a database link in the target database to the clone database:

```
SQL> create public database link apr12_clone connect to sales identified by salespw  
using apr12cln;
```

4. Select from the table in the clone database and use the database link to recreate the dropped table in the target database:

```
SQL> create table europe_sales as select * from europe_sales@apr12_clone;
```

To Recover a Corrupted or Dropped Table by Using RMAN Cataloged Snapshot Copies

In this scenario, a table is dropped and must be recreated from an existing online backup by using RMAN recovery commands.

1. To restore a table from the backup, use SnapCenter on the Oracle host to mount the data file backup. To perform point-in-time recovery, the logs must be present in the active file system. If logs are already pruned, mount the log backups to the Oracle host.

Backup Name	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_14.05.44.2010_1	Log	4/23/2018 2:06:11 PM	Not Applicable	False	Cataloged	10547036
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_14.05.44.2010_0	Data	4/23/2018 2:05:58 PM	Verified	False	Cataloged	10547032
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_13.24.45.7137_1	Log	4/23/2018 1:25:13 PM	Not Applicable	False	Cataloged	10545288
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_13.24.45.7137_0	Data	4/23/2018 1:25:00 PM	Verified	False	Cataloged	10545284
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_08.13.12.3079_1	Log	4/23/2018 8:13:46 AM	Not Applicable	False	Cataloged	10533442
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_08.13.12.3079_0	Data	4/23/2018 8:13:27 AM	Verified	False	Cataloged	10533438
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_08.11.51.8234_1	Log	4/23/2018 8:12:33 AM	Not Applicable	False	Not Cataloged	10533167
RHEL3_demo_netapp_com_SCMPROD_RHEL3_04-23-2018_08.11.51.8234_0	Data	4/23/2018 8:12:18 AM	Unverified	False	Not Cataloged	10533163

2. Select the data file backup that is closest to your requirement and mount it to the Oracle host, then execute the recovery command:

```

RMAN> run {
  recover table ebin.test_restore of pluggable database PDBSCM until time "to_date('04-23-2018 13:25:00','mm/dd/yyyy hh24:mi:ss')" auxiliary destination '/tmp'
}

```

```

Starting recover at 23-APR-18
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=255 device type=DISK
RMAN-05026: WARNING: presuming following set of tablespaces applies to specified Point-in-Time

```

```

List of tablespaces expected to have UNDO segments
Tablespace SYSTEM
Tablespace UNDOTBS1

```

Creating automatic instance, with SID='Bkvo'

initialization parameters used for automatic instance:

```

db_name=SCMPROD
db_unique_name=Bkvo_pitr_PDBSCM_SCMPROD
compatible=12.1.0.2.0
db_block_size=8192
db_files=200
diagnostic_dest=/oracle_home/app
_ system_trig_enabled=FALSE
sga_target=2560M
processes=200
db_create_file_dest=/tmp
log_archive_dest_1='location=/tmp'
enable_pluggable_database=true
clone_one_pdb_recovery=true

```

Performing import of tables...

```

IMPDP> Master table "SYS"."TSPITR_IMP_Bkvo_tigg" successfully loaded/unloaded
IMPDP> Starting "SYS"."TSPITR_IMP_Bkvo_tigg":
IMPDP> Processing object type TABLE_EXPORT/TABLE/TABLE
IMPDP> Processing object type TABLE_EXPORT/TABLE/TABLE_DATA
IMPDP> .. imported "EBIN"."TEST_RESTORE" 5.062 KB 1 rows
IMPDP> Processing object type TABLE_EXPORT/TABLE/STATISTICS/TABLE_STATISTICS
IMPDP> Processing object type TABLE_EXPORT/TABLE/STATISTICS/MARKER
IMPDP> Job "SYS"."TSPITR_IMP_Bkvo_tigg" successfully completed at Mon Apr 23 15:21:48 2018 elapsed 0 00:00:03
Import completed

```

```

Removing automatic instance
Automatic instance removed
auxiliary instance file /tmp/SCMPROD/datafile/o1_mf_temp_ffvyfk0o_.tmp deleted
auxiliary instance file /tmp/SCMPROD/datafile/o1_mf_temp_ffvyfh7g_.tmp deleted
auxiliary instance file /tmp/BKVO_PITR_PDBSCM_SCMPROD/onlineolog/o1_mf_3_ffvyhdm_.log deleted
auxiliary instance file /tmp/BKVO_PITR_PDBSCM_SCMPROD/onlineolog/o1_mf_2_ffvyhdg3_.log deleted
auxiliary instance file /tmp/BKVO_PITR_PDBSCM_SCMPROD/onlineolog/o1_mf_1_ffvyhd84_.log deleted
auxiliary instance file /tmp/BKVO_PITR_PDBSCM_SCMPROD/datafile/o1_mf_users_ffvyh87s_.dbf deleted
auxiliary instance file /tmp/SCMPROD/datafile/o1_mf_sysaux_ffvyf32k_.dbf deleted
auxiliary instance file /tmp/SCMPROD/datafile/o1_mf_system_ffvyf01r_.dbf deleted
auxiliary instance file /tmp/SCMPROD/datafile/o1_mf_sysaux_ffvydk0m_.dbf deleted
auxiliary instance file /tmp/SCMPROD/datafile/o1_mf_undotbs1_ffvydb0h_.dbf deleted
auxiliary instance file /tmp/SCMPROD/datafile/o1_mf_system_ffvyd2v6_.dbf deleted
auxiliary instance file /tmp/SCMPROD/controlfile/o1_mf_ffvycv1m_.ctl deleted
auxiliary instance file tspitr_Bkvo_40368.dmp deleted
Finished recover at 23-APR-18

```

3. After recovering the table, check whether the table exists with exact records:

```

SQL>
SQL> alter session set container=PDBSCM;

Session altered.

SQL>
SQL>
SQL> select * from ebin.test_restore;

      COL1
-----
         1

```

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- NetApp Product Documentation page
<https://docs.netapp.com>
- Microsoft documentation about Active Directory Domains and Trusts
<https://technet.microsoft.com/en-us/library/cc770299.aspx>
- Oracle Databases on ONTAP
<https://www.netapp.com/us/media/tr-3633.pdf>
- NetApp Interoperability Matrix Tool (IMT)
<https://mysupport.netapp.com/matrix/#welcome>
- SnapCenter Software 4.0 Installation and Setup Guide
<http://docs.netapp.com/ocsc-40/topic/com.netapp.doc.ocsc-isg/home.html>
- SnapCenter 4.0 Command Reference Guide
https://library.netapp.com/ecm/ecm_download_file/ECMLP2840882

Version History

Version	Date	Document Version History
Version 1.0	June 2018	Initial Release

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2018 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.