



SnapManager for Oracle uses Protection Manager to protect a database backup

SnapManager Oracle

NetApp
November 04, 2025

Table of Contents

SnapManager for Oracle uses Protection Manager to protect a database backup	1
Details of the target database	1
Primary and secondary storage configuration and topology	1
Backup schedule and retention strategy	5
Workflow summary for local and secondary database backup	6
Protected backup configuration and execution	7
Using SnapManager for Oracle to create the database profile for a local backup	7
Using Protection Manager to configure a secondary resource pool	8
Using Protection Manager to configure secondary backup schedules	9
Using Protection Manager to configure a secondary backup protection policy	10
Using SnapManager for Oracle to create the database profile and assign a protection policy	11
Using Protection Manager to provision the new dataset	13
Using SnapManager for Oracle to create a protected backup	14
Using SnapManager for Oracle to confirm backup protection	14
Database restoration from backup	15
Using SnapManager for Oracle to restore a local backup on primary storage	15
Using SnapManager for Oracle to restore backups from secondary storage	16

SnapManager for Oracle uses Protection Manager to protect a database backup

SnapManager for Oracle and Protection Manager, when installed on a UNIX host and on the server respectively, give the SnapManager database administrator (DBA) the ability to configure and carry out policy-based Oracle database backups to secondary storage, and to restore, if necessary, the backed up data from secondary to primary storage.

In the following example, a DBA, who is using SnapManager, creates a profile for a local backup on primary storage and another profile for a protected backup to secondary storage. Then this DBA works with his network storage administrator, who is using the Protection Manager's console, to configure a policy-based backup of that database from primary to secondary storage.

Details of the target database

This example of integrated database protection describes the protection of a payroll database. The following data is used in the example.

The database administrator (DBA) at TechCo, a 3000-person company headquartered in Atlanta, must create a consistent backup of the production payroll database, PAYDB. The protection strategy for backing up to primary and secondary storage requires that the DBA and the storage administrator work together to back up the Oracle database both locally on primary storage and also remotely, to secondary storage at a remote location.

- Profile information**

When creating a profile in SnapManager, you need the following data:


- Database name: PAYDB
- Host name: payroll.techco.com
- Database ID: payrolldb
- Profile name: payroll_prod
- Connection mode: Database authentication
- Snapshot naming scheme: smo_hostname_dbsid_smoprofile_scope_mode_smid (which translates to "smo_payroll.xyz.com_payrolldb_payroll_prod_f_h_x")

Primary and secondary storage configuration and topology

In this example, the TechCo corporation runs its payroll database on a database server that is also a SnapManager for Oracle host and stores its payroll database data and configuration files on primary storage systems at company headquarters. The corporate requirement is to protect that database with daily and weekly backups to local storage as well as backups to storage systems at a secondary storage site fifty miles away.

The following illustration shows the SnapManager for Oracle and the NetApp Management Console data protection capability components required to support local and secondary backup protection.

Architecture



To manage the payroll database and support its local and secondary backup protection as illustrated in the previous graphic, the following deployment is used.

- **SnapManager host**

The SnapManager host, payroll.techco.com, is located at company headquarters and runs on a UNIX server, which also runs the database program that generates and maintains the payroll database.

- **Connections**

To support local backup and secondary backup protection, the SnapManager host has network connections to the following components:

- SnapManager for Oracle client
- SnapManager repository, which runs the database program, SnapDrive for UNIX, and SnapManager
- Primary storage systems
- Secondary storage systems
- DataFabric Manager server

- **Installed products**

The SnapManager host is installed with the following products for this example:

- SnapManager server

- SnapDrive for UNIX
- Host Utilities

- **TechCo primary storage systems**

The payroll database, including associated data files, log files, and control files, reside on the primary storage systems. These are located at TechCo company headquarters along with the SnapManager host and the network connecting primary storage and the SnapManager host. The latest payroll database transactions and updates are written to the primary storage systems. Snapshot copies, which provide local backup protection of the payroll database, also reside on the primary storage systems.

- **Connections**

To support secondary backup protection, the primary storage systems have network connections to the following components:

- SnapManager host running the database program, SnapDrive for UNIX, and SnapManager
- Secondary storage systems
- DataFabric Manager server

- **Installed products**

The following licenses must be enabled on these systems for this example:

- Data ONTAP 7.3.1 or later
- SnapVaultData ONTAP Primary
- FlexVol (required for NFS)
- SnapRestore
- NFS protocol

- **TechCo secondary storage systems**

The secondary storage systems, located at a network-connected secondary storage site fifty miles away, are used to store secondary backups of the payroll database.

- **Connections**

To support secondary backup protection, the secondary storage systems have network connections to the following components:

- Primary storage systems
- DataFabric Manager server

- **Installed products**

The following licenses must be enabled on the secondary storage systems for this example:

- Data ONTAP
- SnapVaultData ONTAP Secondary
- SnapRestore
- FlexVol (required for NFS)

- NFS protocol

- **DataFabric Manager server**

The DataFabric Manager server, `techco_dfm`, is located at company headquarters in a location accessible by the storage administrator. The DataFabric Manager server, among other functions, coordinates the backup tasks between primary and secondary storage.

- **Connections**

To support secondary backup protection, the DataFabric Manager server maintains network connections to the following components:

- NetApp Management Console
- Primary storage systems
- Secondary storage systems

- **Installed products**

The DataFabric Manager server is licensed for the following server products for this example:

- DataFabric Manager

- **SnapManager repository**

The SnapManager repository, located on a dedicated server, stores data about operations performed by SnapManager, for example the time of backups, tablespaces and datafiles backed up, storage systems used, clones made, and Snapshot copies created. When a DBA attempts a full or partial restore, SnapManager queries the repository to identify backups that were created by SnapManager for Oracle for restoration.

- **Connections**

To support secondary backup protection, the secondary storage systems have network connections to the following components:

- SnapManager host
- SnapManager for Oracle client

- **NetApp Management Console**

The NetApp Management Console is the graphical user interface console used by the storage administrator to configure schedules, policies, datasets, and resource pool assignments to enable backup to secondary storage systems, which are accessible to the storage administrator.

- **Connections**

To support secondary backup protection, NetApp Management Console has network connections to the following components:

- Primary storage systems
- Secondary storage systems
- DataFabric Manager server

- **SnapManager for Oracle client**

The SnapManager for Oracle client is the graphical user interface and command line console used by the DBA for the payroll database in this example to configure and carry out local backup and backup to secondary storage.

- **Connections**

To support local backup and secondary backup protection, SnapManager for Oracle client has network connections to the following components:

- SnapManager host
- SnapManager repository, running the database program, SnapDrive for UNIX, and SnapManager
- Database host (if separate from the host running SnapManager)
- DataFabric Manager server

- **Installed products**

To support local backup and secondary backup protection, the SnapManager for Oracle client software must be installed on this component.

Backup schedule and retention strategy

The DBA wants to ensure that backups are available in case of a loss of data, in case of a disaster, and for regulatory reasons. This requires a carefully thought out retention policy for the various databases.

For the production payroll database, the DBA adheres to the following TechCo retention strategy:

Backup frequency	Retention duration	Backup time	Type of storage
Once daily	10 days	7 p.m.	Primary (local)
Once daily	10 days	7 p.m.	Secondary (archive)
Once weekly	52 weeks	Saturdays 1 a.m.	Secondary (archive)

- **Local backup advantages**

Daily local backup provides database protection, which is instantaneous, uses zero network bandwidth, uses a minimum of additional storage space, provides instantaneous restore, and provides finely-grained backup and restore capability.

Because the final weekly backups of the payroll database are retained for a minimum 52 weeks at a secondary storage site, there is no need to retain the daily backups any longer than 10 days.

- **Protected backup advantages**

Daily and weekly backups to secondary storage at a remote location guarantee that if the data at the primary storage site is damaged, the target database is still protected and can be restored from secondary storage.

The daily backups to secondary storage are made to protect against primary storage system damage.

Because the final weekly backups of the payroll database are retained for a minimum 52 weeks, there is no need to retain the daily backups any longer than 10 days.

Workflow summary for local and secondary database backup

In this example, the DBA (using SnapManager) and the storage administrator (using the NetApp Management Console data protection capability) coordinate actions to configure local backup and secondary backup (also known as a protected backup) of the target database.

The sequence of actions carried out are summarized as follows:

- **Secondary resource pool configuration**

The storage administrator uses the NetApp Management Console data protection capability to configure a resource pool of storage systems at the secondary site that can be used to store the payroll database backup.

- **Secondary backup scheduling**

The storage administrator uses the NetApp Management Console data protection capability to configure secondary backup schedules.

- **Protection policy configuration**

The storage administrator uses the NetApp Management Console data protection capability to configure a secondary backup protection policy for the target database. The protection policy includes the schedules and specifies the base type of protection to implement backup protection (backup, mirror, or a combination of both), and names retention policies for primary data, secondary, and sometimes tertiary storage nodes.

- **Database profile configuration and protection policy assignment**

The DBA uses SnapManager to create or edit a profile of the target database that supports secondary backup. While configuring the profile, the DBA:

- Enables backup protection to secondary storage.
- Assigns the new protection policy, which was created in and retrieved from the NetApp Management Console data protection capability, to this profile.

Assigning the protection policy automatically includes the target database in a partially provisioned, but nonconformant the NetApp Management Console data protection capability dataset. When fully provisioned, the dataset configuration enables backup of the target database to secondary storage.

The dataset name uses this syntax: smo_hostname_databasename, which translates to "smo_payroll.techco.com_paydb".

- **Secondary and tertiary storage provisioning**

The storage administrator uses the NetApp Management Console data protection capability to assign resource pools to provision the secondary and sometimes tertiary storage nodes (if the assigned protection policy specifies tertiary storage nodes).

- **Backup on local storage**

The DBA opens the profile with protection enabled in SnapManager and creates a full backup to local storage. The new backup shows in SnapManager as scheduled for protection, but not yet protected.

- **Secondary backup confirmation**

Because the backup was based on a protection-enabled profile, the backup is transferred to secondary according to the protection policy's schedule. The DBA uses SnapManager to confirm the transferral of the backup to secondary storage. After the backup has been copied to secondary storage, SnapManager changes the backup Protection State from "Not protected" to "Protected."

Protected backup configuration and execution

You must configure SnapManager and Protection Manager to support database backup to secondary storage. The database administrator and the storage administrator must coordinate their actions.

Using SnapManager for Oracle to create the database profile for a local backup

The database administrators use SnapManager to create a database profile that will be used to initiate a backup to local storage on a primary storage system. The entire profile creation and backup creation processes are performed entirely in SnapManager; they do not involve Protection Manager.

A profile contains the information about the database being managed, including its credentials, backup settings, and protection settings for backups. By creating a profile, you do not need to specify database details each time you perform an operation on that database, instead just provide the profile name. A profile can reference only one database. That same database can be referenced by more than one profile.

1. Go to the SnapManager for Oracle client.
2. From the SnapManager Repositories tree, right-click the host you want associated with this profile, and select **Create Profile**.
3. In the Profile Configuration Information page, enter the following information and click **Next**.
 - Profile name: payroll_prod
 - Profile password: payroll123
 - Comment: Production Payroll database
4. In the Database Configuration Information page, enter the following information and click **Next**.
 - Database name: PAYDB
 - Database SID: payrolldb
 - Database host: Accept the default
- Because you are creating a profile from a host in the repository tree, SnapManager displays the host name.
5. In the second Database Configuration Information page, accept the following database information and click **Next**:

- Host Account, representing the Oracle user account: oracle
- Host Group, representing the Oracle group: dba

6. In the Database Connection Information page, select **Use database Authentication** to allow users to authenticate using database information.

For this example, enter the following information and click **Next**.

- SYSDBA Privileged User Name, representing the system database administrator who has administrative privileges: sys
- Password (SYSDBA password): oracle
- Port to connect to database host: 1521

7. In the RMAN Configuration Information page, select **Do not use RMAN** and click **Next**.

Oracle Recovery Manager (RMAN) is an Oracle tool that helps you back up and recover Oracle databases using block-level detection.

8. In the Snapshot Naming Information page, specify a naming convention for the Snapshots associated with this profile by selecting variables. The only variable that is required is the **smid** variable, which creates a unique snapshot identifier.

For this example, do the following:

- a. In the Variable Token list, select the **{usertext}** variable and click **Add**.
- b. Enter "payroll.techco.com_" as the host name and click **OK**.
- c. Click **Left** until the host name appears just after "smo" in the Format box.
- d. Click **Next**.

The Snapshot naming convention of smo_hostname_smoprofile_dbid_scope_mode_smid becomes "smo_payroll.techco.com_payroll_prod2_payrolldb_f_a_x" (where the "f" indicates a full backup, the "a" indicates the automatic mode, and the "x" represents the unique SMID).

9. On the Perform Operation page, verify the information and click **Create**.
10. Click **Operation Details** to see information about the profile create operation and volume-based restore eligibility information.

Using Protection Manager to configure a secondary resource pool

To support backup of the database to secondary storage, the storage administrator uses Protection Manager to organize the secondary storage systems enabled with the SnapVault Secondary license into a resource pool for the backups.

Ideally, storage systems in a resource pool are interchangeable in terms of their acceptability as destinations for backups. For example, when developing the protection strategy for the payroll database, you, as the storage administrator, identified secondary storage systems with similar performance and quality of service levels that would be suitable members of the same resource pool.

You have already created aggregates of unused space on storage systems that you intend to assign to resource pools. This ensures that there is adequate space to contain the backups.

1. Go to Protection Manager's NetApp Management Console.

2. From the menu bar, click **Data > Resource Pools**.

The Resource Pools window appears.

3. Click **Add**.

The Add Resource Pool wizard starts.

4. Complete the steps in the wizard to create the **paydb_backup_resource** resource pool.

Use the following settings:

- Name: Use **paydb-backup_resource**
- Space thresholds (use the defaults):
 - Space utilization thresholds: enabled
 - Nearly Full threshold (for resource pool): 80%
 - Full threshold (for resource pool): 90%

Using Protection Manager to configure secondary backup schedules

To support backup of the database to secondary storage, the storage administrator uses Protection Manager to configure a backup schedule.

Before configuring the schedule for secondary backups, the storage administrator confers with the DBA partner for the following information:

- The schedule that the DBA wants the secondary backups to follow.

In this case, once-daily backups occur at 7 p.m. and once-weekly backups occur on Saturday at 1 a.m.

1. Go to the Protection Manager's NetApp Management Console.
2. From the menu bar, click **Policies > Protection > Schedules**.

The Schedules tab of the Protection Policies window is displayed.

3. Select the Daily schedule **Daily at 8:00 PM** in the list of schedules.
4. Click **Copy**.

A new Daily schedule, **Copy of Daily at 8:00 PM**, is displayed in the list. It is already selected.

5. Click **Edit**.

The Edit Daily Schedule property sheet opens to the Schedule tab.

6. Change the schedule name to **Payroll Daily at 7 PM**, update the description, then click **Apply**.

Your changes are saved.

7. Click the **Daily Events** tab.

The schedule's current Daily backup time of 8:00 p.m. is displayed.

8. Click **Add** and enter **7:00 PM** in the new time field, then click **Apply**.

The schedule's current Daily backup time is now 7:00 p.m.

9. Click **OK** to save your changes and exit the property sheet.

Your new Daily schedule, **Payroll Daily at 7 PM**, is displayed in the list of schedules.

10. Select the Weekly schedule **Sunday at 8:00 PM plus daily** in the list of schedules.

11. Click **Copy**.

A new Weekly schedule, **Copy of Sunday at 8:00 PM plus daily**, is displayed in the list. It is already selected.

12. Click **Edit**.

The Edit Weekly Schedule property sheet opens to the Schedule tab.

13. Change the schedule name to **Payroll Saturday at 1 AM plus daily at 7 PM** and update the description.

14. From the **Daily Schedule** drop-down list, select the Daily schedule you just created, **Payroll Daily at 7 PM**.

Selecting **Payroll Daily at 7 PM** means that this schedule defines when Daily operations occur when the **Payroll Saturday at 1 AM plus daily at 7 PM** schedule is applied to a policy.

15. Click **OK** to save your changes and exit the property sheet.

Your new Weekly schedule, **Payroll Saturday at 1 AM plus daily at 7 PM**, is displayed in the list of schedules.

Using Protection Manager to configure a secondary backup protection policy

After configuring the backup schedule, the storage administrator configures a protected backup storage policy in which that schedule is to be included.

Before configuring the protection policy, the storage administrator confers with the DBA partner for the following information:

- Retention duration to specify for secondary storage
- Type of secondary storage protection required

The protection policy that is created, can be listed in SnapManager for Oracle by the DBA partner and assigned to a database profile for the data to be protected.

1. Go to Protection Manager's NetApp Management Console.

2. From the menu bar, click **Policies > Protection > Overview**.

The Overview tab on the Protection Policies window is displayed.

3. Click **Add Policy** to start the Add Protection Policy wizard.

4. Complete the wizard with the following steps:

- a. Specify a descriptive policy name.

For this example, enter **TechCo Payroll Data: Backup** and a description, then click **Next**.

- b. Select a base policy.

For this example, select **Back up** and click **Next**.

- c. In the Primary Data node policy property sheet, accept the default settings and click **Next**.



In this example, the local backup schedule that was configured in SnapManager is applied. Any local backup schedule that is specified using this method is ignored.

- d. In the Primary Data to Backup connection property sheet, select a backup schedule.

For this example, select **Payroll Saturday at 1 AM plus daily at 7 PM** as your backup schedule, then click **Next**.

In this example, the schedule that you selected includes both the weekly and daily schedules that you configured earlier.

- e. In the Backup policy property sheet, specify the name for the backup node and the retention times for Daily, Weekly, or Monthly backups.

For this example, specify a Daily backup retention of 10 days and a Weekly backup retention of 52 weeks. After you complete each property sheet, click **Next**.

After all property sheets are completed, the Add Protection Policy wizard displays a summary sheet for the protection policy that you want to create.

5. Click **Finish** to save your changes.

The **TechCo Payroll Data: Backup** protection policy is listed among the other policies configured for Protection Manager.

The DBA partner can now use SnapManager for Oracle to list and assign this policy when creating the database profile for the data to be protected.

Using SnapManager for Oracle to create the database profile and assign a protection policy

You must create a profile in SnapManager for Oracle, enable protection in the profile, and assign a protection policy to create a protected backup.

A profile contains information about the database being managed, including its credentials, backup settings, and protection settings for backups. After you create a profile, you do not need to specify database details each time you perform an operation. A profile can reference only one database, but that same database can be referenced by more than one profile.

1. Go to the SnapManager for Oracle client.
2. From the Repositories tree, right-click the host, and select **Create Profile**.
3. On the Profile Configuration Information page, enter the profile details, and click **Next**.

You can enter the following information:

- Profile name: payroll_prod2
- Profile password: payroll123
- Comment: Production Payroll database

4. On the Database Configuration Information pages, enter the database details, and click **Next**.

You can enter the following information:

- Database name: PAYDB
- Database SID: payrolldb
- Database host: Accept the default. Because you are creating a profile from a host in the repository tree, SnapManager displays the host name.
- Host Account, representing the Oracle user account: oracle
- Host Group, representing the Oracle group: dba

5. On the Database Connection Information page, click **Use database Authentication** to allow users to authenticate using database information.

6. Enter the database connection details and click **Next**.

You can enter the following information:

- SYSDBA Privileged User Name, representing the system database administrator who has administrative privileges: sys
- Password (SYSDBA password): oracle
- Port to connect to database host: 1521

7. On the RMAN Configuration Information page, click **Do not use RMAN** and click **Next**.

Oracle Recovery Manager (RMAN) is an Oracle tool that helps you back up and recover Oracle databases using block-level detection.

8. On the Snapshot Naming Information page, specify a naming convention for the Snapshots associated with this profile by selecting variables.

The smid variable creates a unique snapshot identifier.

Perform the following:

- In the Variable Token list, select usertext and click **Add**.
- Enter payroll.techco.com_ as the host name and click **OK**.
- Click **Left** until the host name appears just after smo in the Format box.
- Click **Next**.

The Snapshot naming convention of smo_hostname_smoprofile_dbsid_scope_mode_smid becomes "smo_payroll.techco.com_payroll_prod2_payrolldb_f_a_x" (where "f" indicates a full backup, "a" indicates the automatic mode, and "x" represents the unique SMID).

9. Select **Protection Manager Protection Policy**.

The **Protection Manager Protection Policy** enables you to select a protection policy that was configured

by using NetApp Management Console.

10. Select **TechCo Payroll Data: Backup** as the protection policy from the protection policies retrieved from NetApp Management Console, and click **Next**.
11. On the Perform Operation page, verify the information and click **Create**.
12. Click **Operation Details** to see information about the profile create operation and volume-based restore eligibility information.
 - The assignment of a NetApp Management Console protection policy to the database profile automatically creates a nonconformant dataset, visible to the NetApp Management Console operator, with the name convention `smo_<hostname>_<filename>`, or in this example: `smo_payroll.tech.com_PAYDB`.
 - If the profile is not eligible for volume restore (also called "fast restore"), the following occurs:
 - The **Results** tab indicates that the profile creation was successful and that warnings occurred during the operation.
 - The **Operation Details** tab includes a WARNING log, which states the profile is not eligible for fast restore and explains why.

Using Protection Manager to provision the new dataset

After the `smo_paydb` dataset is created, the storage administrator uses Protection Manager to assign storage system resources to provision the dataset's Backup node.

Before provisioning the newly created dataset, the storage administrator confers with the DBA partner for the name of the dataset specified in the profile.

In this case, the dataset name is `smo_payroll.tech.com_PAYDB`.

1. Go to Protection Manager's NetApp Management Console.
2. From the menu bar, click **Data > Datasets > Overview**.

The Datasets tab of the Datasets window displays a list of datasets that includes the dataset that was just created through SnapManager.

3. Locate and select the `smo_payroll.tech.com_PAYDB` dataset.

When you select this dataset, the graph area displays the `smo_paydb` dataset with its backup node unprovisioned. Its conformance status is flagged as nonconformant.

4. With the `smo_paydb` dataset still highlighted, click **Edit**.

The Protection Manager's NetApp Management Console displays the Edit Dataset window for the `smo_payroll.tech.com_PAYDB` dataset. The window's navigation pane displays configuration options for the dataset's primary node, backup connection, and backup node.

5. From the navigation pane, locate the options for the dataset's backup node and select **provisioning/resource pools**.

The Edit Dataset window displays a setting for default provisioning policy and a list of available resource pools.

6. For this example, select the `paydb_backup_resource` resource pool and click **>**.

The selected resource pool is listed in the "Resource Pools for this node" field.

7. Click **Finish** to save your changes.

The Protection Manager automatically provisions the secondary backup node with resources from the paydb_backup_resource resource pool.

Using SnapManager for Oracle to create a protected backup

When creating a backup for this example, the DBA selects to create a full backup, sets backup options, and selects protection to secondary storage. Although the backup is initially made on local storage, because this backup is based on a protection-enabled profile, the backup is then transferred to secondary storage according to the protection policy's schedule as defined in Protection Manager.

1. Go to the SnapManager for Oracle client.
2. From the SnapManager Repository tree, right-click the profile containing the database that you want to back up, and select **Backup**.

The SnapManager for Oracle Backup Wizard starts.

3. Enter Production_payroll as the label.
4. Enter Production payroll Jan 19 backup as the comment.
5. Select **Auto** as the type of backup that you want to create.

This allows SnapManager to determine whether to perform an online or offline backup.

6. Select **Daily** or **Weekly** as the frequency of the backup.
7. To confirm that the backup is in a valid format for Oracle, check the box next to **Verify backup**.

This operation uses Oracle DBVerify to check the block format and structure.

8. To force the state of the database into the appropriate mode (for example, from open to mounted), select **Allow startup or shutdown of database, if necessary**, and click **Next**.
9. In the Database, Tablespaces, or Datafiles to Backup page, select **Full Backup** and click **Next**.
10. To protect the backup on secondary storage, check **Protect the Backup** and click **Next**.
11. In the Perform Operation page, verify the information you supplied and click **Backup**.
12. In the progress page, view the progress and results of the backup creation.
13. To view the details of the operation, click **Operation Details**.

Using SnapManager for Oracle to confirm backup protection

Using SnapManager for Oracle, you can view a list of backups associated with a profile, determine whether the backups were enabled for protection, and view the retention class (daily or weekly, in this example).

At first, the new backup in this example shows as scheduled for protection, but not yet protected (in the SnapManager graphical user interface and in the backup show command output). After the storage administrator ensures that the backup has been copied to secondary storage, SnapManager changes the

backup protection state from "Not protected" to "Protected" in both the graphical user interface and with the backup list command.

1. Go to the SnapManager for Oracle client.
2. In the SnapManager Repository tree, expand the profile to display its backups.
3. Click the **Backups/Clones** tab.
4. In the Reports pane, select **Backup Details**.
5. View the Protection column and ensure that the status is "Protected."

Database restoration from backup

If the active content of the payroll database is accidentally lost or destroyed, SnapManager and the NetApp Management Console data protection capability support restoration of that data from either a local backup or secondary storage.

Using SnapManager for Oracle to restore a local backup on primary storage

You can restore local backups that exist on primary storage. The entire process is performed using SnapManager for Oracle.

You can also preview information about a backup restore process. You might want to do this to see information about restore eligibility of a backup. SnapManager analyzes data on a backup to determine whether the restore process can be completed by using the volume-based restore or the file-based restore method.

The restore preview shows the following information:

- Which restore mechanism (fast restore, storage-side file system restore, storage-side file restore, or host-side file copy restore) will be used to restore each file.
- Why more efficient mechanisms were not used to restore each file.

In preview of the restore plan, SnapManager does not restore anything. The preview shows information up to 20 files.

If you want to preview a restore of data files but the database is not mounted, then SnapManager mounts the database. If the database cannot be mounted, then the operation fails and SnapManager returns the database to its original state.

1. From the Repository tree, right-click the backup you want to restore, and select **Restore**.
2. On the Restore and Recovery Wizard Welcome page, click **Next**.
3. On the Restore Configuration Information page, select **Complete Datafile/Tablespace Restore with Control Files**.
4. Click **Allow shutdown of database if necessary**.

SnapManager changes the database state, if necessary. For example, if the database is offline and it needs to be online, SnapManager forces it online.

5. On the Recovery Configuration Information page, click **All Logs**.

SnapManager restores and recovers the database to the last transaction and applies all required logs.

6. On the Restore Source Location Configuration page, view the information about the backup on primary and click **Next**.

If the backup exists only on primary storage, SnapManager restores the backup from the primary storage.

7. On the Volume Restore Configuration Information page, select **Attempt volume restore** to attempt volume restore method.

8. Click **Fallback to file-based restore**.

This allows SnapManager to use the file-based restore method if the volume restore method cannot be used.

9. Click **Preview** to see the eligibility checks for fast restore and information about mandatory and overridable checks.

10. On the Perform Operation page, verify the information you have entered, and click **Restore**.

11. To view details about the process, click **Operation Details**.

Using SnapManager for Oracle to restore backups from secondary storage

Administrators can restore protected backups from secondary storage and can choose how they want to copy the data back to the primary storage.

Before you attempt to restore the backup, check the properties of the backup and ensure that the backup is freed on the primary storage system and is protected on secondary storage.

1. From the SnapManager for Oracle Repository tree, right-click the backup you want to restore, and select **Restore**.
2. In the Restore and Recovery Wizard Welcome page, click **Next**.
3. In the Restore Configuration Information page, click **Complete Datafile/Tablespace Restore with Control Files**.
4. Click **Allow shutdown of database if necessary**, and then click **Next**.

SnapManager changes the database state, if necessary. For example, if the database is offline and it needs to be online, SnapManager forces it online.

5. At the Recovery Configuration Information page, click **All Logs**. Then, click **Next**.

SnapManager restores and recovers the database to the last transaction and applies all required logs.

6. In the Restore Source Location Configuration page, select the ID of the protected backup source and click **Next**.
7. In the Volume Restore Configuration Information page, click **Attempt volume restore** to attempt volume restore.
8. Click **Fallback to file-based restore**.

This allows SnapManager to use the file-based restore method if the volume restore method cannot be completed.

9. To see the eligibility checks for fast restore and information about mandatory and overridable checks, click **Preview**.

10. At the Perform Operation page, verify the information you have supplied and click **Restore**.

11. To view details about the process, click **Operation Details**.

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

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

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

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

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

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

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