

15. Backup and restore (in Oracle RDBMS). Monitoring and advisors.

1.1. Oracle recovery functions.

Oracle provides several features that allow you to recover from database failures, including hardware and user errors. A database administrator must be able to protect it against hardware failure by implementing a backup and recovery strategy. Additionally, recovery from user errors is possible using the rollback features.

Instance recovery or crash recovery is a special form of recovery that occurs the first time an Oracle database instance is started after a crash. In crash recovery, the data files are returned to a state consistent with the transactions, preserving all changes committed up to the time the instance failed.

Data file media recovery is the primary form of user-initiated data recovery. It can be used to recover data from a lost or damaged data file, server parameter file (SPFILE), or control file.

Oracle lookback functions support viewing and scrolling data forward and backward in time as follows:

- **Fastback Query:** Allows you to specify a target time and then execute queries to a user database, viewing the results as they would appear at that time.
- **Fastback Versions query:** Allows to see all versions of all rows that have ever existed in one or more tables in a specified time interval.
- **Query Fastback Transaction:** Allows to see the changes made by a single transaction or made by all transactions during a certain period of time.
- **Fastback Table:** Returns a table to its state at a previous point in time.
- **Fastback Drop:** Reverses the effects of a DROP TABLE statement.
- **Fastback Database:** Provides a more efficient alternative to restoring the database at a point in time.

1.1.1. Backup and recovery concepts.

Backups are of two types:

- **Consequently:**
 - It was made in an offline instance;
 - Allows to open the database immediately after a recovery operation;
 - Requires all changes in the redo logs to be applied to the data files.
- **Inconsistent:**

- It is done when the database is open;
- Requires media recovery to occur after file recovery;
- There may be online and archived redo logs that contain changes that have not yet been applied to the data files.

When a database is backed up, copies are made. Data files, control file and archived redo logs (if any). Restoring a database from a backup is simply copying the physical files that make up the database from some backup media (disk or other media) to their locations during normal database operation.

Consistent or inconsistent backups can be made. A sequential backup is one in which the database can be opened immediately after a restore operation. Creating a sequential backup requires that all changes in the redo logs be applied to the data files. The database must be shut down and the instance shut down to perform this type of backup.

With inconsistent online backups, the archived redo logs may contain changes that have not yet been applied to data files. While the database is open, an inconsistent backup may be created. However, to use an inconsistent backup to restore the user database, a media restore must be performed after restoring the database files from backup.

When restoring archived redo logs and backup data files, the Oracle Database server performs a media restore when an attempt is made to open the database. Database transactions in online redo files and archived redo logs that are not yet reflected in the data files are applied to the data files. All pending transactions are cancelled. This puts the data files in a transactionally consistent state before the database is opened.

Media recovery can be either a full recovery or a point-in-time recovery. A full restore applies all changes from the log files and returns the database to its state at the time of failure. It can then be reopened without data loss.

A point-in-time restore restores a database to its state at a specific target time in the past of the user's choosing. The changes are applied to the data files, starting with a set of backup copies of data files created before the target time and a full set of archived redo logs from the time of the backup to the target time. When all changes to the target time are reapplied, the data files are reverted to their contents at the target time. A point-in-time restore is sometimes called a partial restore because not all changes are applied.

A media restore requires a control file, data files that are normally restored from a backup, and all online and archived redo logs from the time the data files

were backed up. Usually it is used only in case of database corruption caused by media corruption, such as file or disk loss.

1.1.2. Media recovery sequence.

Restoring database media includes the following steps (fig. 1.1.):

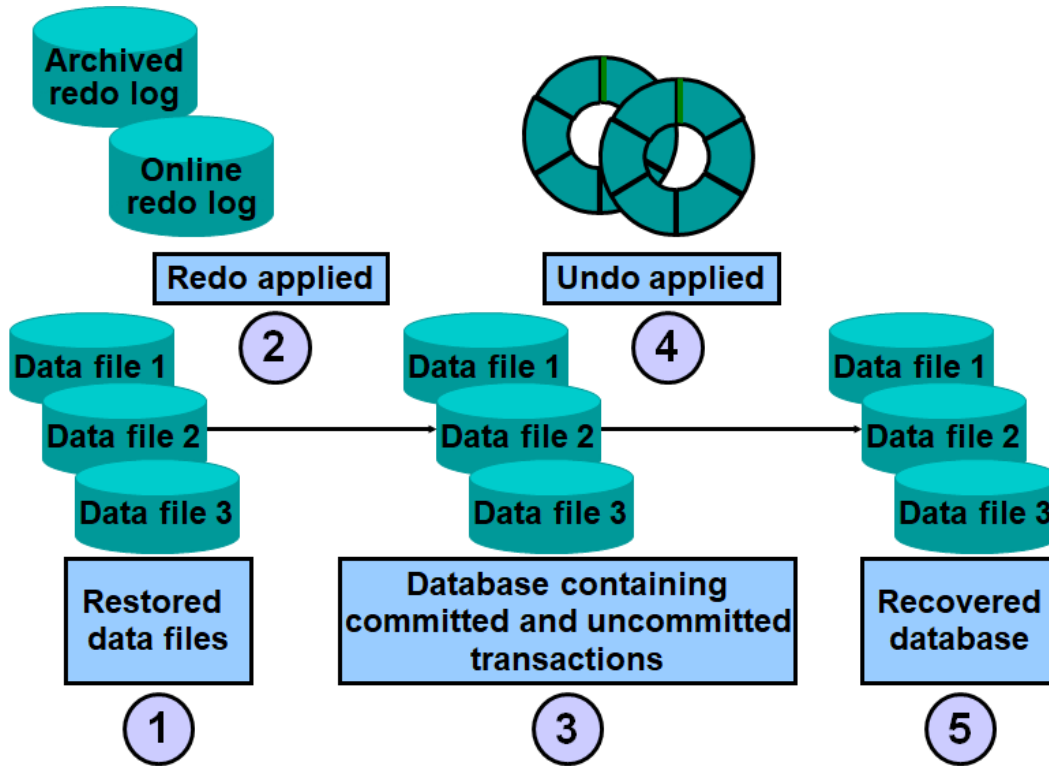


Fig. 1.1. Media recovery sequence.

1. Corrupt or missing files are restored from backup.
2. Changes from archived redo log files and online redo log files are applied as needed. At this point, undo blocks are generated. This is called rolling back or transaction recovery.
3. The database can now contain committed and uncommitted changes.
4. Undo blocks are used to roll back any uncommitted changes. This is known as rolling back or rolling back transactions.
5. The database is now in a restored state.

1.1.3. Configure the database for backup and recovery.

To use Oracle's features to automatically manage backup and restore operations, a database is configured as follows:

- A fast recovery area is used to automate storage management for most backup-related files.

- Work with a given database in ARCHIVELOG mode to be able to do online archiving.

- Uses the fast recovery area as a destination for an archived redo log.

The user can set policies to control which files are backed up, what format is used to store backups on disk, and when files become eligible for deletion from the fast recovery area.

Fast recovery area sizing - the fast recovery area.

Ideally, the fast recovery area should be large enough to hold two full backup copies of user data files, plus any additional backups and backup logs required to restore a user database to any point in time during the custom recovery window.

1.1.4. Configuration of fast recovery area - fast recovery area.

ARCHIVELOG mode and the Fast Recovery Zone are configured to perform this activity:

1. A directory is created to contain the quick recovery area in a given operating system. The user must ensure that the permissions of this directory allow the Oracle Database server to create files in it.

2. Select the Availability property page from the database home page. Then select Recovery Settings in the Backup/Restore Settings area.

The Recovery Settings page appears (fig. 1.2.).



Recovery Settings

Show SQL Revert Apply

Instance Recovery

The fast-start checkpointing feature is enabled by specifying a non-zero desired mean-time to recover (MTTR) value, which will be used to set the FAST_START_MTTR_TARGET initialization parameter. This parameter controls the amount of time the database takes to perform crash recovery for a single instance. When fast-start checkpointing is enabled, Oracle automatically maintains the speed of checkpointing so that the requested MTTR is achieved. Setting the value to 0 will disable this functionality.

Current Estimated Mean Time To Recover (seconds) 12

Desired Mean Time To Recover 0 Minutes

Media Recovery

The database is currently in NOARCHIVELOG mode. In ARCHIVELOG mode, hot backups and recovery to the latest time are possible, but you must provide space for archived redo log files. If you change the database to ARCHIVELOG mode, you should perform a backup immediately. In NOARCHIVELOG mode, only cold backups are possible and data may be lost in the event of database corruption.

ARCHIVELOG Mode*

Log Archive Filename Format* %t_%s_%r.dbf

Number	Archived Redo Log Destination	Status	Type
1	USE_DB_RECOVERY_FILE_DEST	VALID	Local
Add Another Row			

TIP It is recommended that archived redo log files be written to multiple locations spread across the different disks.
 TIP You can specify up to 10 archived redo log destinations.

Enable Minimal Supplemental Logging

Minimal supplemental logging logs the minimal amount of information needed for LogMiner (and any product building on LogMiner technology) to identify, group, and merge the redo operations associated with DML changes.

Fast Recovery

This database is using a fast recovery area. The chart shows space used by each file type

Fig. 1.2. Recovery Settings.

3. Select the ARCHIVELOG mode check box in the media recovery area if it is not already selected. Below the ARCHIVELOG Mode check box is a list of up to 10 possible locations to archive log files. Destination number 1 specifies USE_DB_RECOVERY_FILE_DEST as the destination indicating that the fast recovery area should be used.


4. Specify the name and size of the directory for the fast recovery area in the Fast Recovery Area area on the same page (fig. 1.3.). The location can be a directory in the operating system or the name of an ASM disk group.



The "Apply changes to SPFILE only" checkbox is not selected.

5. Select Apply. A confirmation page appears.

Fast Recovery

This database is using a fast recovery area. The chart shows space used by each file type that is not reclaimable by Oracle. Performing backups to tertiary storage is one way to make space reclaimable. Usable Fast Recovery Area includes free and reclaimable space.

Fast Recovery Area Location 

Fast Recovery Area Size  

Fast Recovery Area Size must be set when the location is set.


Non-reclaimable Fast Recovery Area (B)

Reclaimable Fast Recovery Area (B)

Free Fast Recovery Area (GB)

Enable Flashback Database*

Flashback database can be used for fast database point-in-time recovery, as it returns the database to a prior point-in-time without restoring files. Flashback is the preferred point-in-time recovery method in the recovery wizard when appropriate. The fast recovery area must be set to enable flashback database.

Flashback Retention Time 

Current size of the flashback logs(GB)

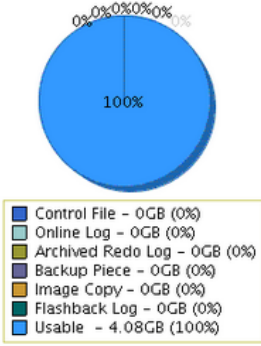
Lowest SCN in the flashback data

Flashback Time

Apply initialization parameter changes to SPFILE only. If not checked, parameter changes will be made to both the SPFILE and the running instance.

* Changes to this setting or parameter require a database restart.

Fast Recovery Area Usage



File Type	Size (GB)	Usage (%)
Control File	0	0%
Online Log	0	0%
Archived Redo Log	0	0%
Backup Piece	0	0%
Image Copy	0	0%
Flashback Log	0	0%
Usable	4.08	100%

Fig. 1.3. Fast recovery settings.

6. Select Yes to restart the instance. This is required to put the database into ARCHIVELOG mode.

7. The Restart Database page: Specify Host and Target Database Credentials page appears. Credentials for the host and databases are provided and confirmed with OK. The Restart Database: Confirmation page appears.

8. It is confirmed (Yes).

The Restart Database: Activity Information page appears. If an error occurs in the browser, the login page will be displayed on refresh.

9. The username and password are set on the database login page.

The Start/Shutdown page appears: Specify host and target database credentials (fig. 1.4.).

10. Provide host and database credentials and then OK.

The database home page appears.

11. It is displayed whether the changes have been implemented.

Restart Database: Specify Host and Target Database Credentials

Specify the following credentials in order to restart the database.

Host Credentials

Specify the OS user name and password to login to target database machine.

* Username

* Password

Database Credentials

Specify the credentials for the target database.
To use OS authentication, leave the user name and password fields blank.

* Username

* Password

Database

* Connect As

Save as Preferred Credential

Note that you need to login to the database as SYSDBA or SYSOPER in order to restart the database.

Fig. 1.4. Instance restart credentials.

1.1.5. Types of archives.

Oracle Enterprise Manager (Enterprise Manager) supports the following archive types:

- Full data file backup: Includes all used blocks of the data file. This can be a backup - an image or a backup copy. Regardless of the format in which the backup is stored, the entire data file is backed up, even if only a few blocks have been changed.
- Incremental data file backups: Includes only those blocks that change between backups in each user database data file. In a typical incremental backup strategy, a level 0 incremental backup that captures all blocks in the data file is taken as the starting point. Subsequent level 1 incremental backups (typically done at regular intervals) capture images of each block in a data file that has changed. Level 1 backups can be cumulative, in which case all blocks changed since the last level 0 backup are included, or they can be differential, in which case only those blocks that have changed since the last level 0 backup are included, or incremental backup level 1.
- Full backups: Includes a backup of the entire database content during the backup. Full backups of all data files are created. Results can be stored as image copies or backup datasets. In both cases, however, the full contents of all database data files are represented in the archive, as well as the control file, the archived redo log, and the server parameter file. With this set of files, the database can be completely restored.

In some circumstances, archiving the entire database is required, such as when switching the database between ARCHIVELOG and NOARCHIVELOG modes.

1.1.6. Archive file types.

With Oracle databases, the following types of archives exist (fig. 1.5.):

- Copies of images. These are copies of a data file, control file, or archived redo log file. A copy can be made using Enterprise Manager or an operating system utility. The image copy of a data file consists of all blocks of the data file, including unused blocks. Copying an image can only include one file. A single copy operation cannot be multiplexed.
- Backup sets - can include one or more data files, a control file, or archived redo log files.

A backup set can be made in two different ways:

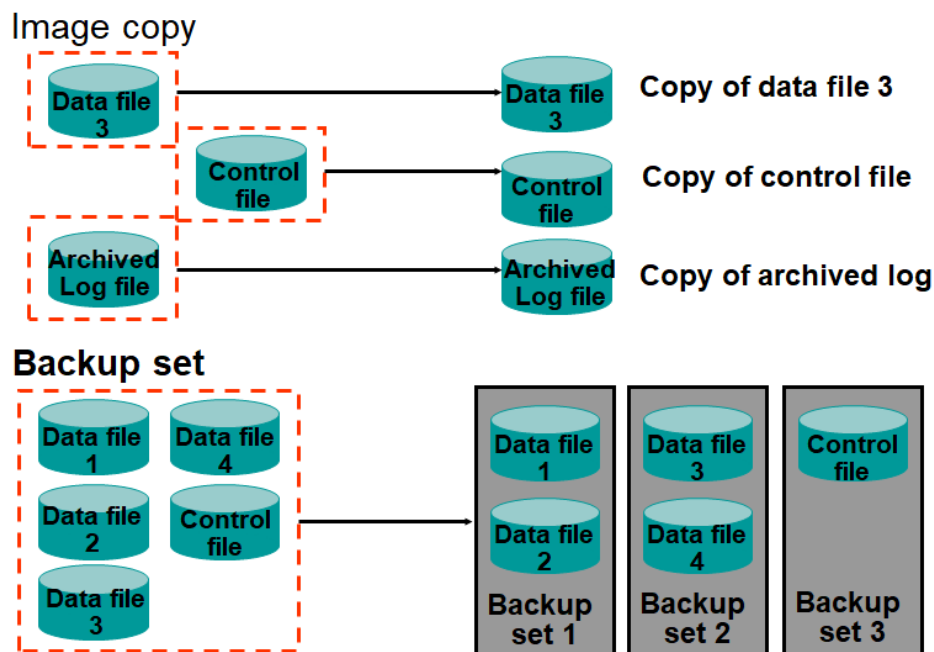


Fig. 1.5. Oracle database backup types.

- Full Backup: A full backup can copy one or more files. In this way, all blocks containing data for the specified files are backed up.
- Incremental backup: An incremental backup is a data file backup that includes only the blocks that have changed since the last incremental backup. Incremental backups require a base level (or incremental level 0) backup that backs up all blocks containing data for the specified files. Incremental level 0 and full

backups copy all blocks in the data files, but full backups cannot be used in an incremental backup strategy.

Automatic control file backup can be configured so that the control file and the current SPFILE are backed up when a BACKUP or COPY command is executed.

The sequentially updated backup feature allows the user to use one or more incremental level 1 backups with an older image backup of the user data files to roll forward the copy to the system change number (SCN) at which the last incremental level 1 backup was made. Any blocks changed since the image copy was created are overwritten with their new contents at the time of the last incremental level 1 backup. The effect is to roll the file forward in time, so that its contents are equivalent to a full backup copy of the data file from an image copy taken during the last incremental level 1 backup. This feature allows strategies with shorter recovery times to be implemented.

1.2. Configuration of the backup settings.

1.2.1. Storage settings.

The user can configure settings that apply to all backup jobs on the Configure Backup Settings page. These settings can be changed if necessary.

Using Enterprise Manager, to configure disk settings, a user have to:

1. Select Backup Settings in the Backup / Recovery area of the Availability page. The Disk Settings section of the Configure Backup Settings page is displayed.
2. The default value of 1 for Parallelism is selected.
3. The disk backup location field should be blank, indicating that the backups will be placed in the Quick Recovery area.
4. Default Backup Set is selected for the disk backup type.
5. Specify the host credentials for the backup in the Host Credentials region.
6. Select the Test Disk Backup button to verify that the backup credentials and location are correct.
7. The disk backup test success confirmation message is displayed.

1.2.2. Backup policy settings.

Backup policies may include:

- SPFILE backup (server settings file);
- archive the control file;
- skipping unchanged files;
- enable change tracking.

Information retention policies may include:

- backup copies to save;
- recovery window;
- truncation / repetition.

The Policy Settings page sets backup policies governing control file and SPFILE backups, specifies tablespaces to be excluded from all database backups, and sets the backup saving policy.

To configure backup policies using Enterprise Manager, following best practices, a user should do the following:

1. Select Configure backup settings in the Backup / Restore area of the Availability page.

2. Select the Policy tab (fig. 1.6.).

3. Check the box "Automatically back up the control file and server parameter file (SPFILE) with every backup and database structural change". No location is specified for "Auto Backup Disk Location" so that backups are made to the Fast Recovery Area.

4. The "Optimize the whole database backup by skipping unchanged files such as offline and read-only data files that have been backed up" box is checked.

5. Check the box "Enable block change tracking for faster incremental backups". If no database area is configured to use server-managed files, then these files will be kept in the database area (fig. 1.6.).

6. Select "Retain backups that are necessary for a recovery to any time within the specified number of days (point-in-time recovery)" - by default 31 days (fig. 1.6.).

7. Select Delete archive logs after specifying the number of archives (default is 1 archive).

8. Specify the host credentials for the backup in the Host Credentials region.

8. It is confirmed - OK.

Next is the home page of the user database.

Backup Policy

Automatically backup the control file and server parameter file (SPFILE) with every backup and database structural change

Autobackup Disk Location

An existing directory or diskgroup name where the control file and server parameter file will be backed up. If you do not specify a location, the files will be backed up to the fast recovery area location.

Optimize the whole database backup by skipping unchanged files such as read-only and offline datafiles that have been backed up

Enable block change tracking for faster incremental backups

Block Change Tracking File ←

Specify a location and file, otherwise an Oracle managed file will be created in the database area.

Tablespaces Excluded From Whole Database Backup

Populate this table with the tablespaces you want to exclude from a whole database backup. Use the Add button to add tablespaces to this table.

Select	Tablespace Name	Tablespace Number	Status	Contents
<input type="checkbox"/>	No Items Selected			

TIP These tablespaces can be backed up separately using tablespace backup.

Retention Policy

Retain All Backups
You must manually delete any backups

Retain backups that are necessary for a recovery to any time within the specified number of days (point-in-time recovery)

Days
Recovery Window

Retain at least the specified number of full backups for each datafile

Backups
Redundancy

Fig. 1.6. Backup policies.

1.2.3. Backing up the entire database.

When a full backup of the user database is created, the whole database is backed up. Full backups of all data files are created. Results can be stored as image copies or backup sets. In both cases, however, the full contents of all database data files are represented in the archive, as well as the control file, the archived redo log, and the SPFILE. With this set of files, the database can be completely restored.

Archiving the whole database is necessary in some circumstances, such as when switching the database between ARCHIVELOG and NOARCHIVELOG modes.

Backing up the database.

To perform a full backup using Enterprise Manager:

1. Select Backup Schedule in the Backup/Recovery area of the Availability page.

The Schedule Backup page appears (fig. 1.7.).

2. Select Whole database from the list of objects available for backup.

Schedule Backup

Oracle provides an automated backup strategy based on your disk and/or tape configuration. Alternatively, you can implement your own customized backup strategy.

Oracle-Suggested Backup

Schedule a backup using Oracle's automated backup strategy.

[Schedule Oracle-Suggested Backup](#)

This option will back up the entire database. The database will be backed up on daily and weekly intervals.

Customized Backup

Select the object(s) you want to back up.

[Schedule Customized Backup](#)



Whole Database

You may only perform an offline backup of the entire database. If the database is OPEN at the time of backup, the database will be shut down and mounted before the backup. The database will be opened after the backup.



All Recovery Files on Disk

Includes all archived logs and disk backups that are not already backed up to tape.

Backup Strategies

Oracle-suggested:

- Provides an out-of-the-box backup strategy based on the backup destination
- Sets up recovery window for backup management
- Schedules recurring and immediate backups
- Automates backup management

Customized:

- Specify the objects to be backed up
- Choose disk or tape backup destination
- Override the default backup settings
- Schedule the backup

Host Credentials

To perform a backup, supply operating system login credentials to access the target database.

* Username

oracle

* Password


••••••

Fig. 1.7. Backup scheduling.

3. The Host Credentials login information is entered. "Schedule Customized Backup" is selected. The page with options of specific schedule backup is displayed (fig. 1.8.).

4. Select Full Backup in the Backup Type area.

5. Select Next - Settings.



Schedule Customized Backup: Options

Database **orcl** [Cancel](#) [Step 1 of 4](#) [Next](#)

Backup Strategy **Customized Backup**

Object Type **Whole Database**

Backup Type

Full Backup

Use as the base of an incremental backup strategy

Incremental Backup

A level 1 cumulative incremental backup includes all blocks changed since the most recent level 0 backup.

Refresh the latest datafile copy on disk to the current time using the incremental backup

Advanced

Delete obsolete backups

Delete backups that are no longer required to satisfy the retention policy.

Use proxy copy supported by media management software to perform a backup

If proxy copy of the selected files is not supported, a conventional backup will be performed.

Maximum Files per Backup Set

Section Size KB

Backs up large files in parallel, using sections of the specified size. (This parameter overrides Maximum Backup Piece Size in Backup Settings.)

[Encryption](#)

[Return to Schedule Backup](#) [Cancel](#) [Step 1 of 4](#) [Next](#)

Fig. 1.8. Schedule customized backup.

6. Select Disk as the backup destination. Moving forward - schedule.
7. Select the default job name and edit the Job Description field as needed. Selects the default start time of "Once (immediate)" to run the backup immediately. Alternatively, one can schedule the job at a specific time in the future and set the parameters as needed.
8. Next step - Next. The Backup Scheduling: Overview page is displayed (fig. 1.9.).
9. Final changes can be made on the Scheduled Backup: Overview page. Then select Submit Job to run the backup job with the specified options. Offline backup is running.

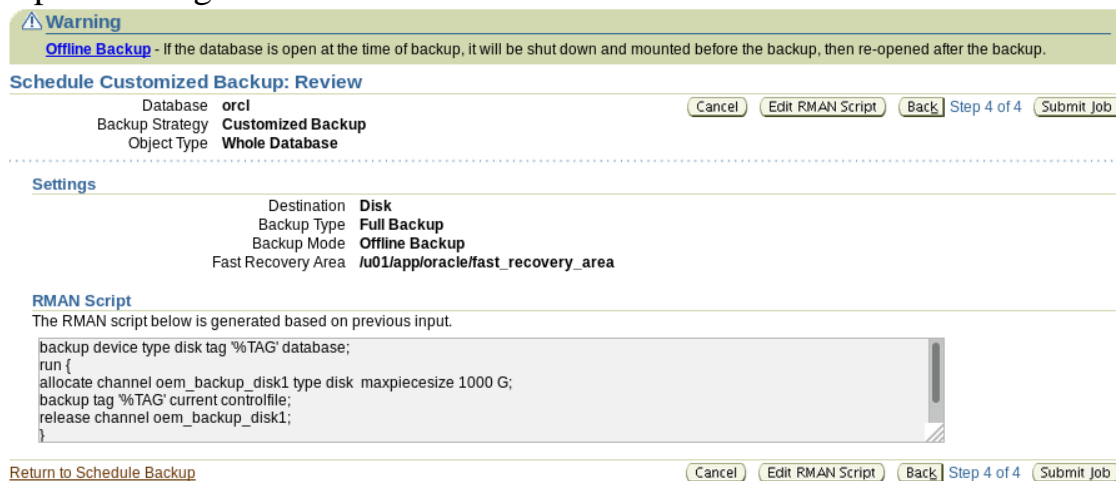


Fig. 1.9. Schedule customized backup: Review.

10. Select View Job to monitor the progress of the backup job. A page is displayed that shows a summary of the recorded job (fig. 1.10.). In the Logs region, user can track the progress of the various steps of the backup job and move down to see the logs of Recovery Manager (RMAN) work on the jobs.

Summary

Status **Succeeded**
 Scheduled **Nov 1, 2021 5:00:00 PM (UTC+02:00)**
 Started **Nov 1, 2021 5:00:00 PM (UTC+02:00)**
 Ended **Nov 1, 2021 5:00:36 PM (UTC+02:00)**
 Elapsed Time **36 seconds**
 Notification **No**

Type **Database Backup**
 Owner **SYS**
 Description **Whole Database Backup**
 Oracle Home [/u01/app/oracle/product/11.2.0/d...](#)
 Oracle SID **orcl**
 Host Username **oracle**
 Database Username **SYS**
 Database Role *********
 Backup Strategy **advanced**
 Version 10g or higher **YES**
 Database Connect String [\(DESCRIPTION=\(ADDRESS_LIST=\(ADDR...](#)
 Database Name **ORCL**
 Blackout **NO**
 Encryption Mode **None**
 Offline Backup **YES**
 Backup Script [Show](#)

Targets

Status **All**

[Expand All](#) | [Collapse All](#)

Name	Targets	Status	Started	Ended	Elapsed Time (seconds)
Execution: orcl	orcl	Succeeded	Nov 1, 2021 5:00:00 PM (UTC+02:00)	Nov 1, 2021 5:00:36 PM (UTC+02:00)	36
Step: Prebackup	orcl	Succeeded	Nov 1, 2021 5:00:16 PM (UTC+02:00)	Nov 1, 2021 5:00:17 PM (UTC+02:00)	1
Step: Backup	orcl	Succeeded	Nov 1, 2021 5:00:26 PM (UTC+02:00)	Nov 1, 2021 5:00:27 PM (UTC+02:00)	1
Step: Post Backup	orcl	Succeeded	Nov 1, 2021 5:00:36 PM (UTC+02:00)	Nov 1, 2021 5:00:36 PM (UTC+02:00)	0

Fig. 1.10. Summary of recorded backup job.

1.3. Backup strategy suggested by Oracle.

Enterprise Manager makes it easy to set up backups with an Oracle-proposed backup strategy that protects user data and provides an efficient recovery capability. The strategy proposed by Oracle uses Oracle's sequential backup capabilities and updated backup features to provide faster recovery than is possible by applying database changes from the backup log to the user data files. The backup strategy depends on the backup devices the user plans to use.

The backup strategy proposed by Oracle is based on creating an image copy of the user database that is rolled forward using constantly updated backups. Enterprise Manager schedules RMAN backup jobs to run nightly.

For each data file, the strategy requires a backup to be made as follows:

- At the beginning of day 1 of the strategy (the time when the first scheduled job actually runs), an incremental backup of a level 0 data file is created. It contains the contents of the data file at the beginning of day 1. In a situation of boot and restore, day 1 redo logs can be used to restore to any point during day 1.

- At the start of day 2, an incremental level 1 backup containing the blocks changed during day 1 is created. In a boot and restore situation, this incremental level 1 can be applied to quickly restore the backup created from level 0 to the start of day 2, and redo logs can be used to restore to any point during day 2.

- At the beginning of each day n for day 3 and later, the level 1 backup from the beginning of day $n-1$ is applied to a level 0 backup. This brings the data file copy to its state at the beginning of day $n-1$. A new level 1 is then created containing the blocks changed on day $n-1$. In a boot and restore situation, this incremental level 1 can be applied to quickly restore a restored backup to the beginning of day n , and the redo logs can be used to restore the database to any point in day n .

The copies of data files that are used in the Oracle-suggested backup strategy are marked with the ORA\$OEM_LEVEL_0 tag. The sequential level 1 backups used in this strategy are designed for use with copies of data files that are labeled accordingly. The user can safely implement other backup strategies without worrying about conflicts with the backups of the strategy proposed by Oracle.

1.3.1. Application of the strategy suggested by Oracle.

We are about to follow how a full backup is done this way using Enterprise Manager:

1. Select Backup Schedule in the Backup/Recovery area of the Availability page. The Schedule Backup page appears.
2. Enter the credentials in the appropriate credentials section.

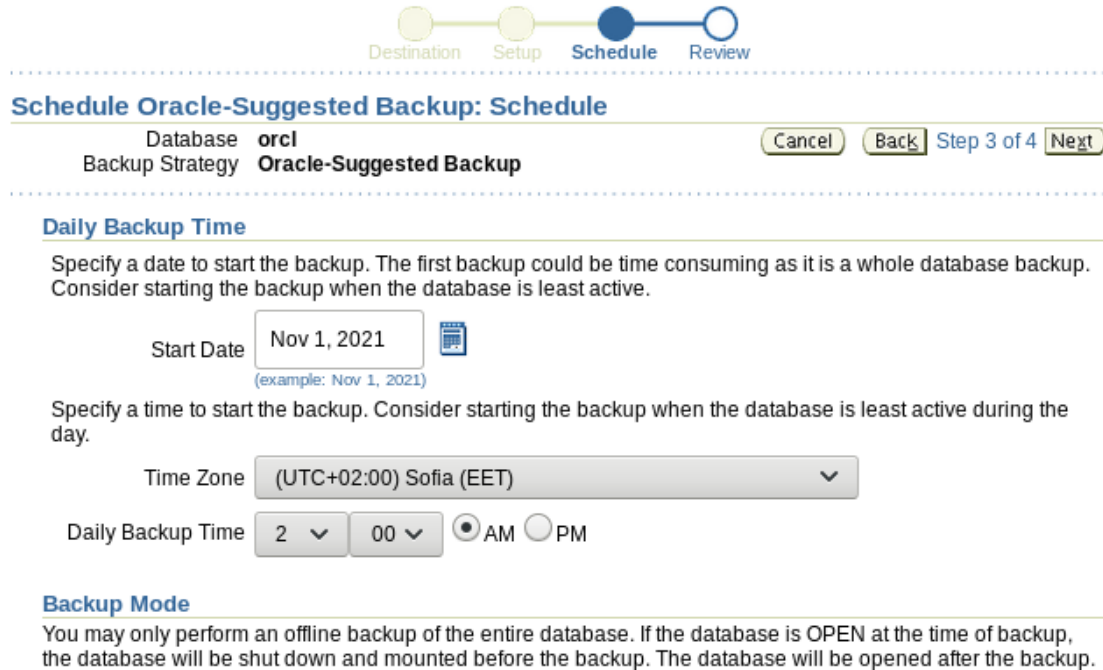
3. Select Schedule Oracle-Suggested Backup. The Schedule Oracle-Suggested Backup: Destination page appears.

4. Select the backup destination Disk. Then - Next. The Schedule Oracle-Suggested Backup: Setup page appears.

5. Check the settings and then continue - Next.

The Oracle Suggested Scheduling: Action Schedule page appears.

6. Date and time of backup is indicated. To the next step - Next. The page for planning the backup over time appears (fig. 1.11.).



Destination Setup **Schedule** Review


Schedule Oracle-Suggested Backup: Schedule

Database **orcl** Cancel Back Step 3 of 4 Next


Backup Strategy **Oracle-Suggested Backup**



Daily Backup Time

Specify a date to start the backup. The first backup could be time consuming as it is a whole database backup. Consider starting the backup when the database is least active.

Start Date 
(example: Nov 1, 2021)

Specify a time to start the backup. Consider starting the backup when the database is least active during the day.

Time Zone 

Daily Backup Time   AM PM

Backup Mode

You may only perform an offline backup of the entire database. If the database is OPEN at the time of backup, the database will be shut down and mounted before the backup. The database will be opened after the backup.

Fig. 1.11. Scheduling backups over time.

7. Settings are being reviewed. Select Submit Job to schedule the backup strategy suggested by Oracle. The Status page confirms the recording of the job.

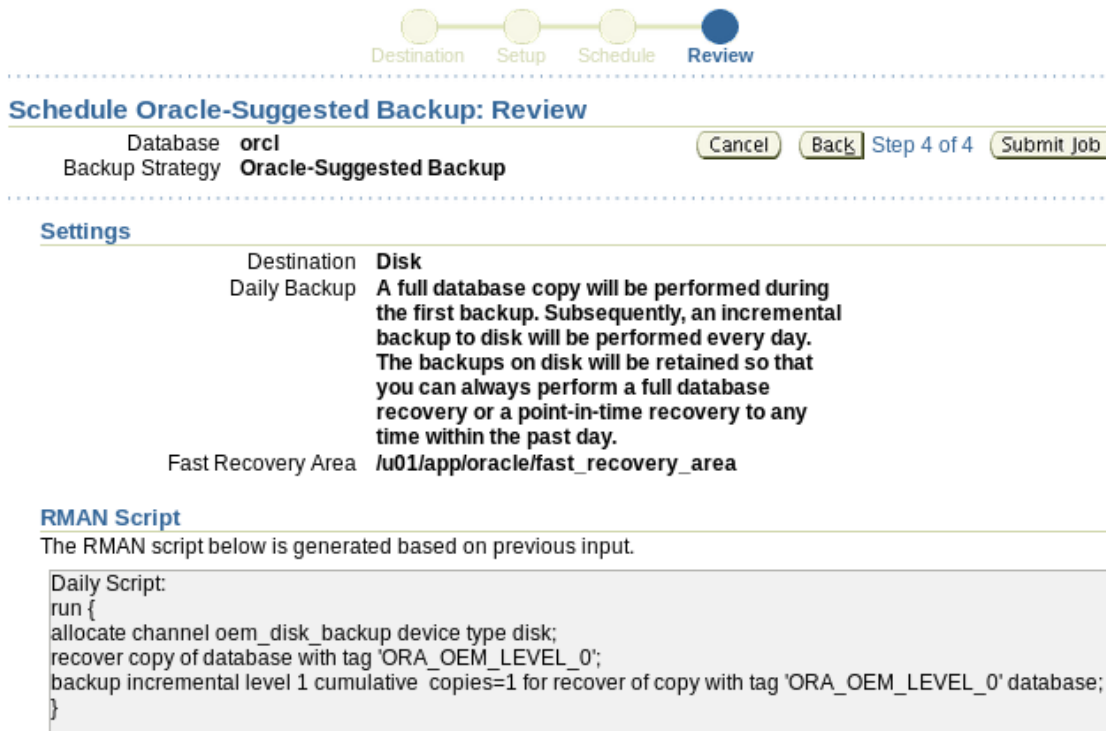


Fig. 1.12. Backup review.

8. Again, from the View Job button, the sent job is reviewed (fig. 1.12.), or from OK, the user returns to the Availability page.

1.4. Backup management.

- User archive management includes:
 - management of archive files on disk or other media;
 - management of archive records.
- Archives can be in one of three states:
 - available: the archive is still present on disk or other media (as recorded in the repository);
 - expired: the archive has been deleted from disk or other media, but is still listed in the repository.
 - unavailable: the archive is temporarily unavailable for data recovery operations.

Backup management consists of managing the backups themselves (as they exist on disk or tape) and managing the recording of backups. Backup records are stored in RMAN (Oracle Recovery Manager) storage.

User backup records can be managed through Enterprise Manager. Backup support features provided in Enterprise Manager include the following:

- View lists of backups (backup defines image copies) stored in the RMAN repository.

- Repository Cross Validation:

- checking if the archives listed in the repository exist and are accessible;
- marking as expired all backups that are not available during the crosscheck.

- Delete records of expired backups from the RMAN repository.

- Delete old archives from storage and from disk.

If a fast recovery zone is used for backup storage, many maintenance activities are reduced or eliminated due to automatic disk space management and fast recovery zone file retention based on retention policy.

1.4.1. Manage Current Backups Page.

In Enterprise Manager, the Manage Current Backups page can be accessed on the Availability page. The Manage Current Backups page has two property pages: Backup Sets (the initial view) and Image Copies (fig. 1.13.). Each serves a similar purpose, listing the backups written to the RMAN repository.

Archive sets are identified by their labels and completion times. To view information about which files are backed up in a backup set, specify the value in the Tag column. For information about the individual spare parts in the spare set, select the number of components in the Pieces column. The individual pieces will be sorted by file name.

The Image Copies property page provides information similar to the Backup Sets property page.

Manage Current Backups

This backup data was retrieved from the database control file.

Catalog Additional Files Crosscheck All Delete All Obsolete Delete All Expired

Backup Sets Image Copies

Search

Status: Available

Contents: Datafile Archived Redo Log SPFILE Control File

Completion Time: Within a month

Go

Results

Crosscheck Change to Unavailable Delete Validate

Select All Select None

Select	Key	Tag	Completion Time	Contents	Device Type	Status	Keep	Pieces
<input type="checkbox"/>	9	TAG20070717T213251	Jul 17, 2007 9:33:01 PM	SPFILE, CONTROLFILE	DISK	AVAILABLE	NO	1
<input type="checkbox"/>	8	TAG20070712T210315	Jul 12, 2007 9:03:23 PM	SPFILE, CONTROLFILE	DISK	AVAILABLE	NO	1
<input type="checkbox"/>	3	BACKUP_ORCL_000001_071007101627	Jul 10, 2007 10:19:14 PM	ARCHIVED LOG	DISK	AVAILABLE	NO	1
<input type="checkbox"/>	1	BACKUP_ORCL_000001_071007101627	Jul 10, 2007 10:18:39 PM	DATAFILE	DISK	AVAILABLE	NO	1

Fig. 1.13. Current archives managements.

1.4.2. Perform backup maintenance tasks.

- Crosschecking: Checking that the actual physical state of the backup matches the backup record in the RMAN repository.
- Deleting expired backups (Deleting expired backups): Deleting backups that RMAN found to be unavailable during a crosscheck operation.
- Deleting obsolete backups (Deleting obsolete backups): Deleting backups that are no longer needed based on the retention policy.

Cross check.

Backup cross-checking causes RMAN to verify that the actual physical state of the backup matches the backup entry in the RMAN repository (fig. 1.14.). During the crosscheck operation, the repository is updated to reflect the current state of the archive. Disk backups are marked as AVAILABLE if they are still present on disk in the location specified in the RMAN repository and if they do not have file header corruption. Backups on other media are listed as AVAILABLE if they are still on other media. (files are not checked for corruption here.) Backups that are missing or corrupted are marked as EXPIRED.

Selecting Crosscheck All at the top of the Manage Current Backups page crosschecks all files in the RMAN repository.

Crosscheck All: Specify Job Parameters

An Enterprise Manager job will be created to perform the specified operation on all backup sets and image copies. Please specify the parameters to run the job.

* Job Name

* Job Description

Schedule

Type One Time (Immediately) One Time (Later) Repeating

Fig. 1.14. Crosschecking parameters.

Individual files can also be cross-checked by selecting the file in the results list and then specifying cross-check at the top of the results list. Unlike Crosscheck All, crosschecking of individual files is done immediately.

Delete expired backups.

When an expired backup is deleted from the RMAN repository, the records of those backups that are marked as expired are deleted.

Expired backups can be deleted with Enterprise Manager by selecting Delete All Expired at the top of the Manage Current Backups page. The special thing here

is that this will delete both expired backup sets and expired image copies from the RMAN repository, regardless of whether the Backup Sets or Image Copies property page is viewed when Delete All Expired is selected.

When Delete All Expired is selected, the Delete All Expired: Specify Job Parameters page is displayed (fig. 1.15.). To check that RMAN has the most up-to-date information, select "Perform "Crosscheck All" Before "Delete All Expires"".

Delete All Expired: Specify Job Parameters

Cancel Show RMAN Script **Submit Job**

An Enterprise Manager job will be created to perform the specified operation on all backup sets and image copies. Please specify the parameters to run the job.

* Job Name Bkp_Mgmt_orcl_000143

* Job Description Backup Management Job for Delete All Expired

Perform the operation 'Crosscheck All' before 'Delete All Expires'.
*Crosscheck All' will update the latest status of the backup sets and image copies.

Schedule

Type One Time (Immediately) One Time (Later) Repeating

Fig. 1.15. Delete all expired archives.

Deleting Obsolete Backups.

Obsolete backups are backups that are no longer needed based on retention policy. Outdated archives can be deleted by selecting Delete All Outdated at the top of the Manage Current Archives page. All obsolete backups (both backup sets and image copies) will be deleted, whether while viewing the Backup Set or Image Copy property page.

When Deleting all obsoletes, the Delete all obsoletes page appears: Specify job parameters (fig. 1.15.). The delete job can be run immediately or scheduled as with a backup job.

If a fast recovery area is used as the only destination for disk-based backups, there will be no need to delete outdated backups from disk. Files will be managed as specified in the archive retention policy and will only be deleted when space is needed.

Marking backups as "unavailable".

If an individual backup is known to be unavailable due to a temporary condition, such as a disk drive being temporarily offline or other media (tape) stored off-site, the backups can be marked as "unavailable". RMAN will retain the backup

information in the RMAN repository (and will not delete it when expired backups are deleted), but will not attempt to use the backup in restore operations. When the backup becomes available again, its status can be changed back to "available".

A backup can be marked as "unavailable" by checking the checkbox next to each archive in the backup results list and selecting "Change Unavailable".

It is Unable to mark backups stored in flash recovery area as unavailable.

Cataloging of additional archives.

If backups are available in the fast recovery area, or backups are made using operating system commands, they can be cataloged in the RMAN repository so that RMAN can use them in a restore operation.

Backups can be added to the catalog by selecting the Catalog Additional Files page at the top of the Manage Current Backups page. On the page, you can choose either "Catalog all files in the recovery area in the Recovery Manager repository" or "Catalog files in the specified disk location in the Recovery Manager repository" (fig. 1.16.).

Catalog Additional Files

Catalog all files in the recovery area into the Recovery Manager repository

Catalog files in the specified disk location into the Recovery Manager repository

Starts With

Enter the directory and the first few characters of the filename, e.g. /usr/oracle/dest.

Fig. 1.16. Cataloging additional files.

1.4.3. Perform an Oracle Advised Recovery.

The Data Recovery Wizard is an Oracle database feature that automatically diagnoses data corruption, determines and presents appropriate repair options, and performs repairs upon user request. Oracle Advised Recovery uses the Data Recovery Advisor.

Errors in the database can be detected by a *data integrity check*. A data integrity check is a diagnostic procedure performed by Health Monitor to assess the health of the database or its components. Data integrity checks are invoked reactively when an error occurs. Checks can also be performed manually.

When a fault is diagnosed, it is recorded in the Automatic Diagnostic Repository (ADR). The data recovery wizard can be used to generate repair advice and repair failures after the failure is detected and stored in the ADR.

The Data Recovery Wizard can be used to use repair options such as:

- blocking media recovery;
- data file media recovery;
- Oracle Flashback Database.

Typically, the data recovery wizard presents options for automated and manual data repair.

Usage of Oracle Advised Recovery to recover the entire database:

1. The Availability page with properties is opened. To call this page, select Perform Recovery.

2. The failure information is reviewed in the recovery section recommended by Oracle. "Advise and Recover" is selected.

The *View and Manage Failures* page appears.

3. The fault description can be expanded for additional information. Select the failure and then - Advice.

The Recovery Advice page is displayed, listing the RMAN script that will be used to recover the data file.

4. Select Continue.

5. Review the information and then select Submit Recovery Job. The Recovery Results page indicates that the restore was successful.

6. Select Open Database.

1.4.4. Load a table with a previous version of data - Flashback Table.

• Flashback Table allows to restore a table (or tables) to a specific point in time without restoring a backup. When this function is used, the data in the tables and all associated objects (indexes, constraints, triggers, etc.) are restored.

• Data is retrieved from the undo tablespace to perform a Flashback Table operation.

• The Flashback Table privilege on a table is required to flash back a table.

• Row movement must be enabled for the table on which the Flashback operation is being performed.

The Flashback Versions and Flashback Transaction queries can be used to determine the appropriate flashback time. Flashback Table provides a way for users to easily and quickly recover from accidental modifications without DBA

involvement. The FLASHBACK TABLE or FLASHBACK ANY TABLE system privileges must be granted to any user who will use the Flashback Table feature. In addition, SELECT, INSERT, DELETE, and ALTER rights must be granted to the user.

Enterprise Manager Can be used to "flash back" a table. The wizard guides the user through the process.

1.4.5. Enable movement of table rows.

In order for a table to be flashed back in time, table row movement must be enabled. When row movement is enabled, the Oracle server can now move a row in the table.

Using Enterprise Manager to enable row movement on a table:

1. Select Tables in the Database Objects area of the Schema property page.
2. Enter the schema name to search the table, then OK.
3. Select the table and Edit.

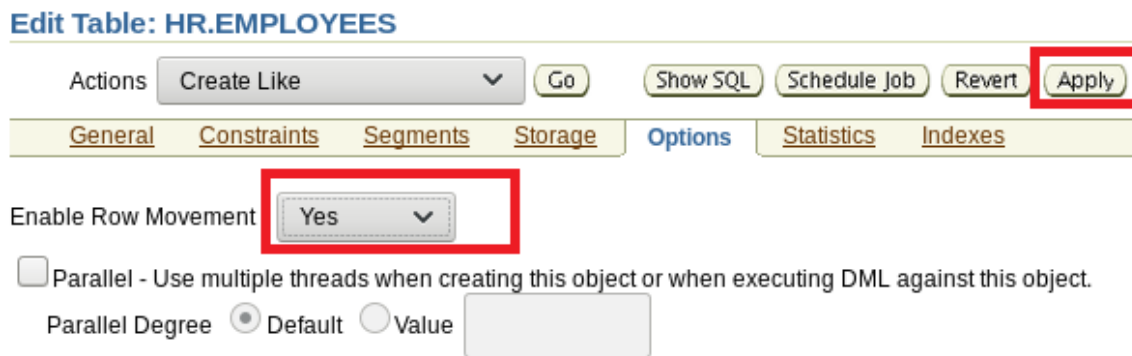


Fig. 1.17. Edit a table.

The Edit Table page appears (fig. 1.17.).

4. Select the Options tab.

5. In the Allow Row Movement drop-down menu, select Yes and Apply to update the options for the table.

The update confirmation message is displayed.

1.4.6. Performing a Flashback Table.

Flash back a table using Enterprise Manager:

1. Select Perform Recovery in the Backup/Restore region of the Availability property page. The page appears.

2. In the User Directed Recovery region, select Tables from the Recovery Scope drop-down list.

The page has been updated.

3. Select Flashback Existing Tables as the operation type and then Recover (fig. 1.18.).

The “Perform Object-Level Recovery: Point-in-time Options” page is displayed.



Fig. 1.18. Performing a Flashback Table.

4. Select “Flashback to a timestamp” and specify the time stamp to which the data will be restored (fig. 1.19.).

The user can also select "Evaluate row and transactions to decide on a point in time" and specify a table name, or one can select "Flashback to a known SCN (system change number)" and to specify the SCN.

Point-in-time Flashback Versions Query Filter Choose SCN Flashback Tables Dependency Options Dependencies More


Perform Object Level Recovery: Point-in-time

Recovery Scope **Tables** Cancel Step 1 of 7 **Next**


Operation Type **Flashback Existing Tables**

Specify the point in time to which to recover.


Evaluate row changes and transactions to decide on a point in time

* Table 
Example: SCOTT.EMP

Flashback to a timestamp

Date  Time AM PM
Example: Mar 19, 2003

Flashback to a restore point

Restore Point 

Flashback to a known SCN

SCN

Fig. 1.19. Timestamp to which the data will be restored.

5. Select Add Tables.

The Flashback Tables: Add tables page appears (fig. 1.20.).

6. Enter the schema and then Search for a list of tables.

7. Select the tables to be returned and then OK.

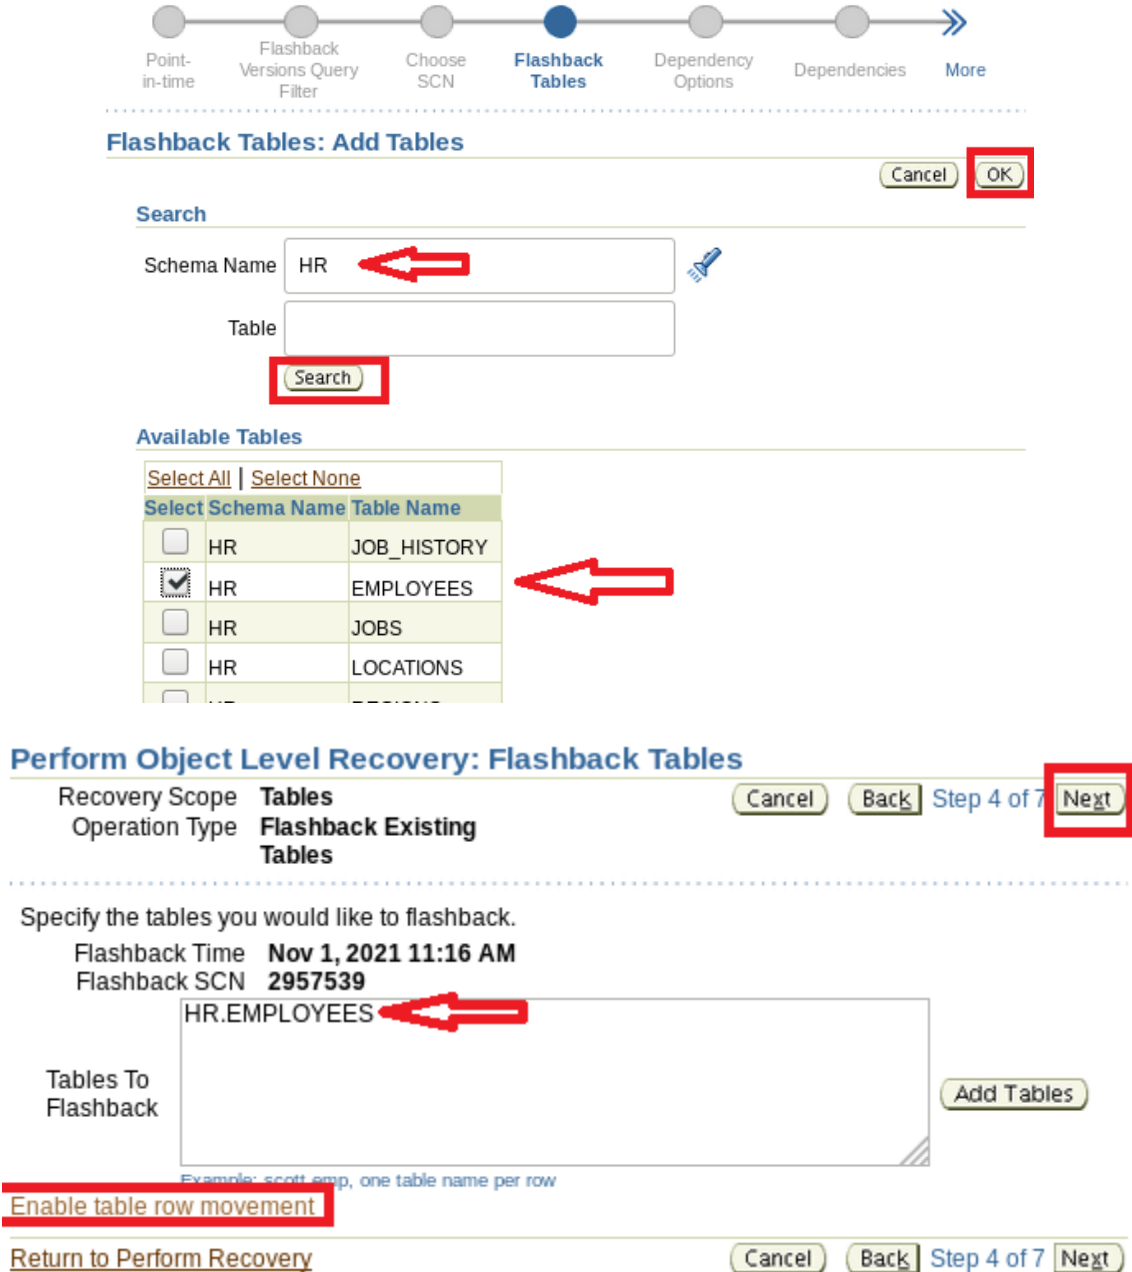


Fig. 1.20. Flashback Tables: adding tables.

If there are any dependencies, the Dependency Options page appears.
8. It is chosen to review the dependencies by selecting the corresponding button (fig. 1.21.).

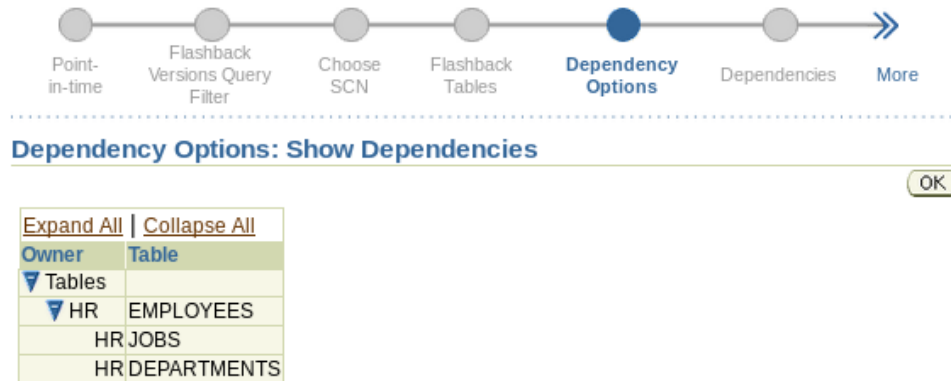


Fig. 1.21. Flashback Tables: Show Dependencies.

9. Select Cascade to flash back the selected tables and all dependent tables. One can also choose one of the following:

- Customize: to flash back the selected tables and some of the dependent tables. This can lead to inconsistent data.
- Restrict: to flash back only the selected tables. This can also lead to inconsistent data.

10. Select Next to continue. The preview page shows which tables will be flashed back.

11. If everything is OK, select Submit to execute Flashback of the table.

1.4.7. Recovering Tables by usage of Flashback Drop.

- Flashback Drop allows to restore a deleted table, returning the deleted table and its dependent objects to the database.
- Deleted tables are placed in the trash until they are purged from there.
- Flashback Drop restores tables from the Recycle Bin.

Using Enterprise Manager, for Flashback Drop:

1. Select Perform Recovery in the Backup/Restore region of the Availability property page. The page appears.
2. In the restore region, the user selects Tables from the Restore Scope dropdown list. The page has been updated.
3. Select Flashback Dropped Tables as the operation type. Then - Recovery. The Perform object-level recovery: Selection of deleted objects page is displayed (fig. 1.22.).
4. The name of the scheme and the name of the table are indicated similarly. The flashlight icon can be selected to search for the schematic. Go is selected.

The page refreshes and the deleted table and its dependent objects are displayed in the Results region.

5. Select the table and Next.

●
○
○

Dropped Objects Selection Rename Review

Perform Object Level Recovery: Dropped Objects Selection

Recovery Scope **Tables** Cancel Step 1 of 3 Next

Operation Type **Flashback Dropped Tables**

Select the tables from the Recycle Bin that you would like to recover. The Results table shows dependent objects that will also be recovered when the selected tables are recovered.

Search

Schema Name Table Go

Results

[Select All](#) | [Select None](#) | [Expand All](#) | [Collapse All](#)

Select	Object Name	Schema	Recovery Scope	Tablespace	Drop Time	Create Time	Size	Operation
<input type="checkbox"/>	Recycle Bin							View Content
<input checked="" type="checkbox"/>	123	HR	TABLE	USERS	2021-10-30:22:45:21	2021-10-30:22:28:23	0	View Content

Fig. 1.22. Perform object-level recovery: select deleted objects.

The Restore: Rename page is displayed.

6. The table can be given a new name or the default name can be selected, the Perform Object-Level Restore: Overview page is displayed.

7. Confirm the request, Submit, to perform the recovery operation. The confirmation page is displayed.

8. Confirm with OK.

2. Monitoring databases and using advisors. Monitoring and advisors.

2.1. Database monitoring.

An Oracle DBA should be familiar with Oracle's self-monitoring architecture and be able to use performance advisors to optimize database performance.

2.1.1. Proactive database monitoring.

Using Oracle Enterprise Manager, one can proactively monitor the health and performance of a given database. By proactively monitoring certain metrics – such as where the Oracle Database server is spending CPU time or how disk space is being used each hour or day, the user can take the necessary corrective steps to avoid future performance issues.

Proactive monitoring includes the following tasks:

- monitoring of the general condition of the database and workload;
- performance monitoring;
- use of alerts.

Oracle Database includes a self-diagnostic engine called the Automatic Database Diagnostic Monitor (ADDM). ADDM enables the Oracle Database server to diagnose its performance and determine how identified problems can be resolved.

2.1.2. Monitoring the overall health of the database and workload.

The status of a database can be monitored using the Enterprise Manager database home page. This page provides general information about the status of the database and reports information that is useful for monitoring the health and load of the database. It is updated periodically.

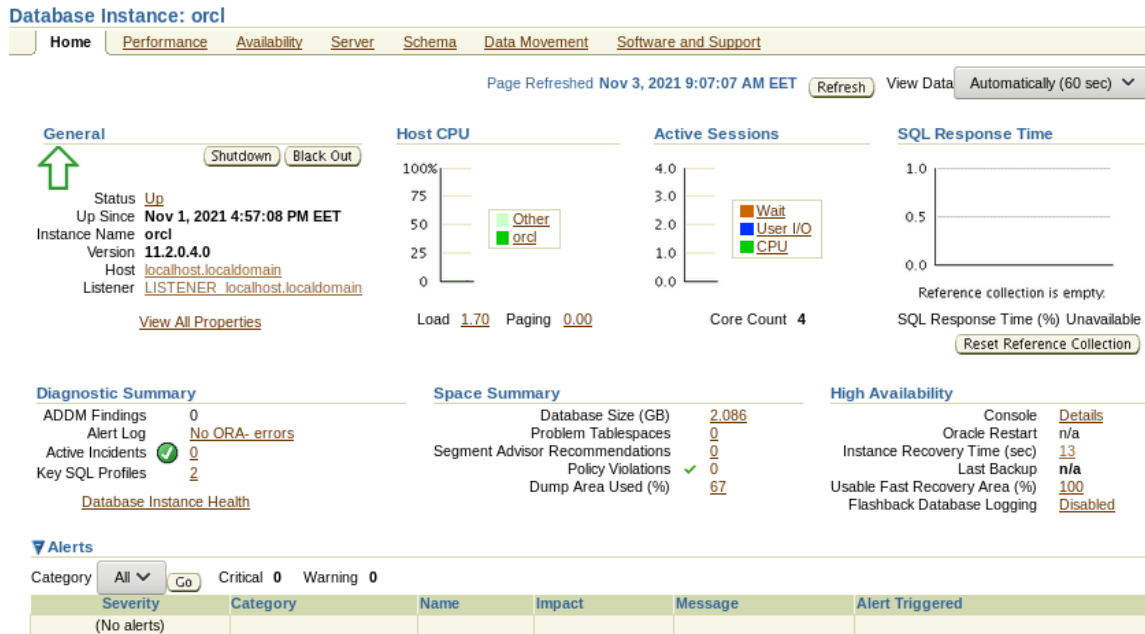


Fig. 2.1. The home page of the database.

The home page of the database (Fig. 2.1.), includes the following:

- **General region (General):** Provides a quick overview of the database, including the status of the database instance, the time the database instance was started, the name of the instance, and the name of the computer system (host).
- **Region for the processor (Host CPU):** shows the percentage of CPU time used throughout the system. This chart divides the CPU percentage into time used by the database instance and time used by other processes. If the user instance of the database is taking up most of the CPU time, the reason can be investigated in detail by looking at the summary of active sessions.
- **Alerts table:** Provides information on all alerts issued, as well as the severity of each. An alert is a notification that a metric threshold has been exceeded.

2.2. Performance monitoring.

To see the overall state of the user system over time in a graphical format (including CPU usage, memory usage and disk usage) the Performance page is used (Fig. 2.2.). The information is graphed over time to help identify periods of time or increased activity. This page covers three main performance areas: host, sessions, and instance throughput. When selecting on the graphs, detailed information is obtained.

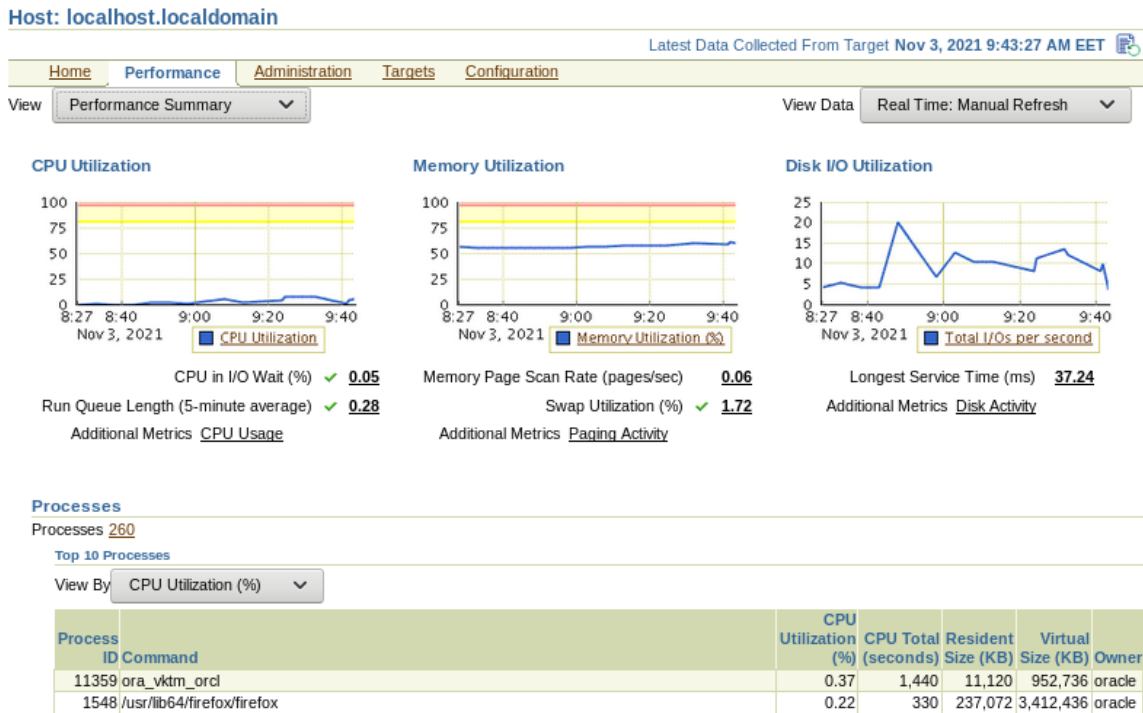


Fig. 2.2. Page Performance.

The top ten processes on the CPU can be seen on this page. Automatic Database Diagnostic Monitor (ADDM) and Active Session History Report (ASH) are usually started.

2.2.1. Using Alerts.

Alerts help proactively monitor a given database. Most alerts are notifications that are generated when certain metric thresholds are exceeded. Critical and warning threshold values can be set for each alert. These threshold values are intended to be threshold values that, when exceeded, indicate that the system is in an undesirable state. In addition to the notification, alerts can be set to perform an action (such as running a script).

The following alerts are enabled by default:

- tablespace usage (warning at 85% full; critical at 97% full);
- the snapshot is too old;
- recovery area with little free space;
- the resuming session is stopped.

When any warning is received, the advisory that is provided is followed, or ADDM or other advisor may be run, as appropriate, to obtain detailed diagnostics of the system or object behavior.

2.2.2. View metrics and thresholds.

Metrics are a set of statistics about certain system attributes as defined by Oracle Database. They are calculated and stored in the Autoload Repository (AWR) and displayed on the All Metrics page (Fig. 2.3.). The All Metrics page can be accessed by selecting the All Metrics link under the Related Links heading on the database home page.

When a link to a specific metric is selected, a details page displays more information about the metric. Online help for this page provides a description of the metric.

For each of these indicators, warning and critical threshold values can be defined. Whenever the threshold is exceeded, the Oracle Database server issues a warning. Alerts are displayed on the database home page under the heading Alerts (or under Related Alerts for non-database alerts).

All Metrics Collected From Target Nov 3, 2021 1:00:30 PM EET

[Expand All](#) | [Collapse All](#)

Metrics	Thresholds	Collection Schedule	Upload Interval	Last Upload
▼ orcl				
▶ Archive Area	Some	Every 15 Minutes	Every Collection	Nov 3, 2021 12:49:45 PM
▶ Data Failure	Some	Every 5 Minutes	Every Collection	-
▶ Database Files	None	Server Generated	Server Generated	Server Generated
▶ Database Job Status	All	Every 5 Minutes	On Alert	Oct 2, 2021 1:30:48 PM
▶ Database Limits	Some	Server Generated	Server Generated	Server Generated
▶ Database Replay	None	Real-time Only	n/a	n/a
▶ Database Replay Client	None	Server Generated	Server Generated	-
▶ Database Services	None	Server Generated	Server Generated	Server Generated
▶ Database Vault Attempted Violations - Command Rules	None	Every 1 Hour	On Alert	-
▶ Database Vault Attempted Violations - Realms	None	Every 1 Hour	On Alert	-
▶ Database Vault Configuration Issues - Command Rules	All	Every 1 Hour	Every Collection	-
▶ Database Vault Configuration Issues - Realms	All	Every 1 Hour	Every Collection	-
▶ Database Vault Policy Changes	All	Every 1 Hour	Every Collection	-
▶ Deferred Transactions	All	Every 5 Minutes	On Alert	Oct 2, 2021 1:31:59 PM
▶ Dump Area	Some	Every 15 Minutes	Every Collection	Nov 3, 2021 12:49:45 PM
▶ Efficiency	None	Server Generated	Server Generated	Server Generated
▶ Failed Logins	All	Every 30 Minutes	Every Collection	-
▶ Fast Recovery	None	Every 15 Minutes	Every Collection	Nov 3, 2021 1:00:18 PM
▶ Health Check	Some	Every 15 Seconds	On Alert	Nov 1, 2021 4:57:33 PM
▶ Incident	Some	Every 5 Minutes	Every Collection	-
▶ Incident Status	All	Every 5 Minutes	Every Collection	Nov 3, 2021 12:58:04 PM
▶ Invalid Objects	None	Every 24 Hours	On Alert	-

Fig. 2.3. Page All Metrics.

2.2.3. Set metric thresholds.

Oracle Server provides a set of predefined metrics, some of which initially have thresholds defined for them.

Existing threshold settings can be modified and threshold values set for other metrics by following these steps:

1. Selects the Metrics and Policy Settings link from the Related links panel on the database home page. The Indicator Thresholds page is displayed (Fig. 2.4.).
2. In the View list, select Thresholded Metrics to view only those thresholded metrics that are either predefined by Oracle or predefined by the user.
3. To set or modify a warning threshold (Warning Threshold) for a specific indicator, enter the desired value in the Warning Threshold field for this indicator.
4. To set or change a critical threshold for a specific indicator, enter the desired value in the field Critical threshold for this indicator.
5. The collection schedule of the metric can be changed by selecting the link in the Collection Schedule field and entering the appropriate information.
6. Select the single pencil icon to use the Edit Advanced Settings page to make changes to Corrective Actions (monitoring), template overrides, and advanced threshold settings.
7. Select the triple pencil icon to set different threshold values for different instances of the object type being measured.
8. Confirm the changes with OK.


▶ User Block	All	Server Generated	Server Generated	Server Generated
▶ User Block Chain	None	Disabled	On Alert	-
▶ Wait Bottlenecks	None	Server Generated	Server Generated	Server Generated
▶ Waits by Wait Class	None	Server Generated	Server Generated	Server Generated

Related Links

Metric and Policy Settings	User-Defined Metrics
--	--------------------------------------

Metric and Policy Settings Cancel OK

Metric Thresholds Policies

View Metrics with thresholds 





Metric	Comparison Operator	Warning Threshold	Critical Threshold	Corrective Actions	Collection Schedule	Edit
Access Violation	Matches	<input type="text"/>	.*	None	Every 5 Minutes	
Access Violation Status	>	<input type="text"/>	0	None	Every 5 Minutes	
Archive Area Used (%)	>	80	<input type="text"/>	None	Every 15 Minutes	
Archiver Hung	Matches	<input type="text"/>	.*	None	Every 5 Minutes	

Fig. 2.4. Metrics and Policy Settings.

Metrics are important for measuring database health and serve as input for self-tuning and recommendations made by Oracle Advisors.

2.2.4. Set up a direct alert notification.

Raised (current) alerts are displayed on the Enterprise Manager database home page. Optionally, Oracle Enterprise Manager Database Control can be directed to provide notifications when events occur that require user intervention.

Setting up an email notification is done by following these steps:

1. Select the Setup link in the header or footer area of a database management page.
2. Select Notification methods on the Setup page (Fig. 2.5.).

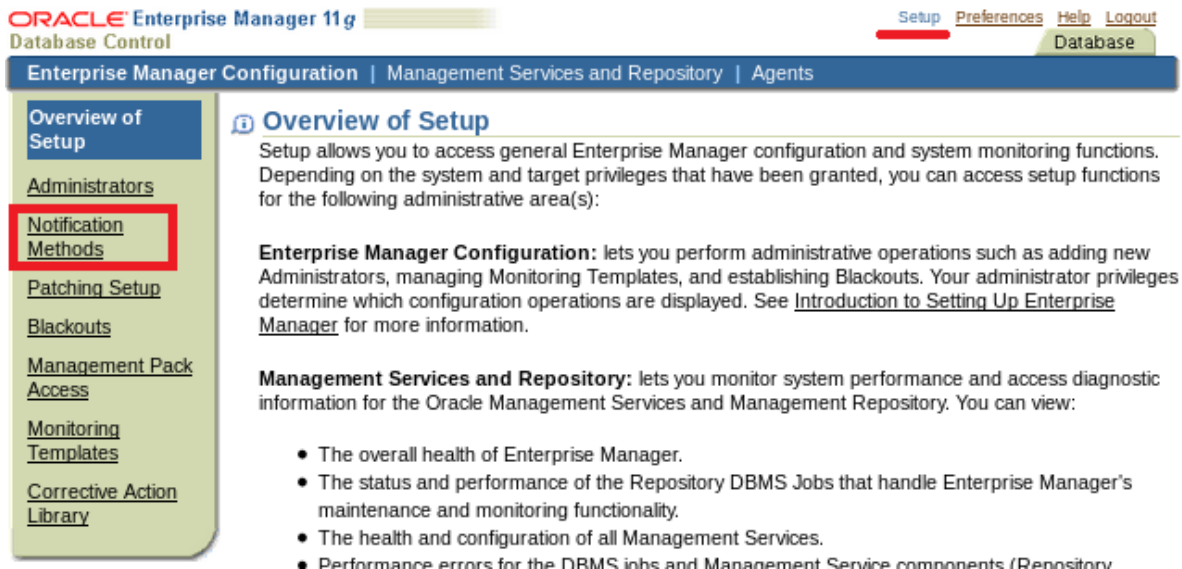


Fig. 2.5. Notification methods.

3. Enter the required information in the Mail Server region on the Notification Methods page. A notification method is set and notification rules are established, an email address is provided to receive notifications.

4. Select the Preferences link in the header or footer area of any database control page.

5. Select General on the Preferences page (Fig. 2.6.). Select Add another row in the email address region to add a new email address.

6. Specify an email address and Apply.

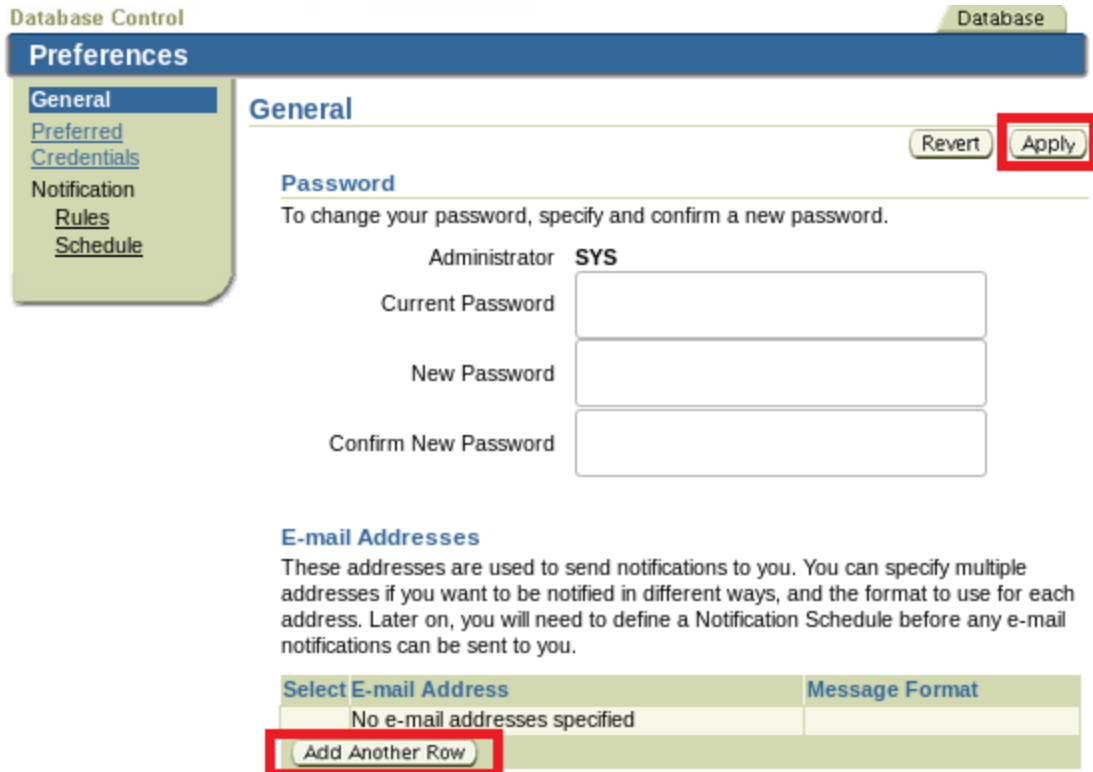


Fig. 2.6. Notification Preferences.

7. Select Rules in the notification area.
8. Select the conditions for which a notification must be made. Apply (Fig. 2.7.).

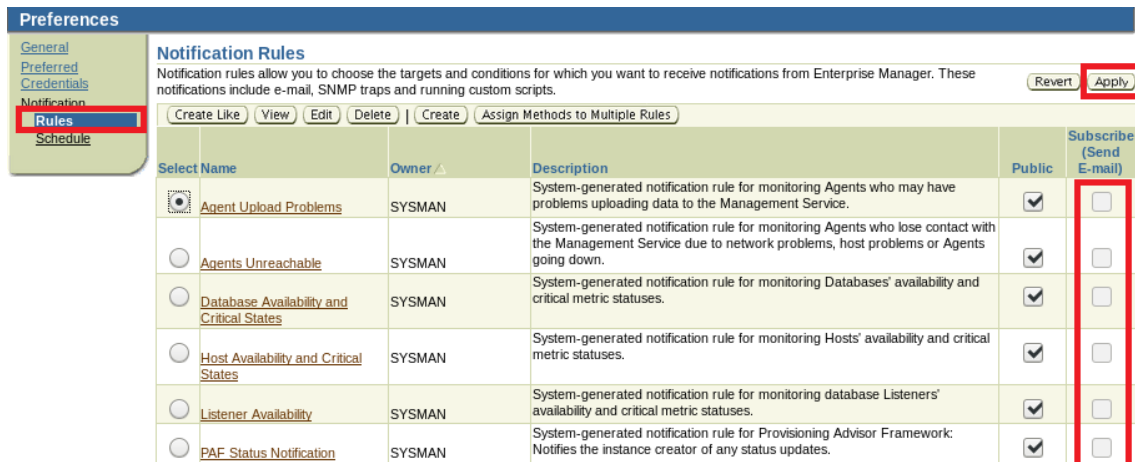


Fig. 2.7. Rules in the notification area.

2.3. Diagnose performance issues.

Database performance issues are flagged by ADDM. To facilitate automatic performance diagnostics using ADDM, the Oracle server periodically collects

information about the state of the database and the workload being monitored. This information is collected in the form of snapshots, which are a statistical summary of the state of the system at any particular point in time. These snapshots are stored in the Automatic Workload Repository (AWR) located in the SYSAUX tablespace. Snapshots are stored in AWR for a period of time and then purged to make space for new snapshots. ADDM inspects the data stored in AWR and performs proactive analysis to determine underlying issues in the database.

ADDM can greatly reduce the amount of effort required to diagnose and tune the performance of Oracle-based systems. ADDM is the general database wizard. It concentrates on those components and operations that consume maximum database time.

In most systems, the majority of performance problems are related to "bottlenecks" that cause significant shortages or overuse of resources. ADDM can identify these bottlenecks in the Oracle server using active session history and snapshot statistics captured in AWR. Working with other server management components and wizards, ADDM can either troubleshoot a problem or make recommendations about available troubleshooting options.

A summary of the results of the Automatic Database Diagnostic Monitor (ADDM) operation is usually reviewed.

On the database control home page, the Diagnostic Summary tab can be used, which gives the number of ADDM findings from the previous autorun.

Selecting the link listed in ADDM findings opens the Automatic Diagnostic Database Monitor (ADDM) page where details of the last ADDM run can be accessed (Fig. 2.8.).

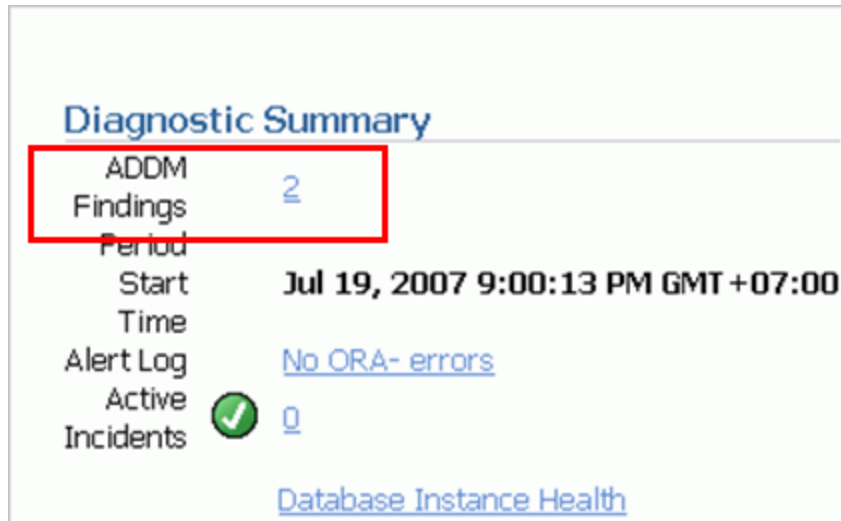


Fig. 2.8. Automatic Database Diagnostic Monitor (ADDM).

2.3.1. View performance analysis information.

ADDM runs automatically every 60 minutes to coincide with the pictures taken by AWR. Its output consists of a description of any problem that is identified that cannot be fixed automatically, as well as a recommended course of action.

In the ADDM Performance Analysis section of the database home page, the cause of the performance issue can be viewed by selecting the find link. The performance findings detail page then describes the findings and recommended actions (Fig. 2.9.).

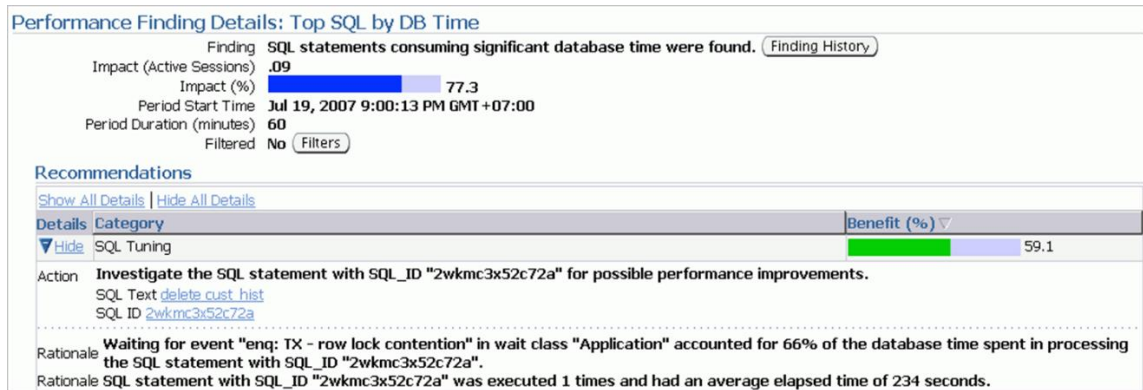


Fig. 2.9. Page with performance findings and recommended actions.

2.3.2. Responding to the findings.

To respond to a performance finding, that finding is selected and the recommended actions, if any, are followed. The recommendation may include the work of an Advisor. Selecting the Find link takes you to the Find Performance details page.

A finding to respond to is selected based on the percentage benefit and then the recommended action is followed. The recommendation may include running an advisor, which can be done by selecting Run Advisor Now.

2.3.3. Change the frequency and retention of AWR snapshots.

The behavior and analysis of ADDM is based on the Automatic Workload Repository (AWR), which collects system performance statistics and stores the data in the database. After a default installation, AWR captures data every 60 minutes and purges data that is more than seven days old. It is possible to configure both the snapshot frequency and the data retention period. For example, the snapshot interval can be shortened when debugging errors described in the ADDM result.

The following settings can be viewed and changed on the autoloading repository page:

- snapshot retention period (initially set to seven days);
- snapshot collection interval (default and recommended value: 60 minutes).

The settings are changed by following these steps:

1. Select Automatic Workload Repository in the statistics management region of the server page (Fig. 2.10).

Database Instance: orcl

Home Performance Availability **Server** Schema Data Movement

Storage
[Control Files](#)
[Tablespaces](#)
[Temporary Tablespace](#)
[Groups](#)
[Datafiles](#)
[Rollback Segments](#)
[Redo Log Groups](#)
[Archive Logs](#)
[Migrate to ASM](#)
[Make Tablespace Locally Managed](#)

Database Configuration
[Memory Advisors](#)
[Automatic Undo Management](#)
[Initialization Parameters](#)
[View Database Feature Usage](#)

Oracle Scheduler
[Jobs](#)
[Chains](#)
[Schedules](#)
[Programs](#)
[Job Classes](#)
[Windows](#)
[Window Groups](#)
[Global Attributes](#)
[Automated Maintenance Tasks](#)

Statistics Management
Automatic Workload Repository
[AWR Baselines](#)

Resource Manager
[Getting Started](#)
[Consumer Groups](#)
[Consumer Group](#)
[Mappings](#)
[Plans](#)

Security
[Users](#)
[Roles](#)
[Profiles](#)
[Audit Settings](#)
[Transparent Data](#)

Fig. 2.10. Automatic Workload Repository.

2. Select Edit on the Automatic Repository for Workload page (Fig. 2.11.). The Edit Settings page appears.
3. A new state snapshot retention period or a new system snapshot interval is introduced. Confirm with OK.

Automatic Workload Repository

Page Refreshed Nov 3, 2021 9:27:25 PM EET

The Automatic Workload Repository is used for storing database statistics that are used for performance tuning.

General

Snapshot Retention (days) **8**
Snapshot Interval (minutes) **60**
Collection Level **TYPICAL**
Next Snapshot Capture Time **Nov 3, 2021 10:00:51 PM**

Manage Snapshots and Baselines

Snapshots **75**
Baselines **1**
Latest Snapshot Time **Nov 3, 2021 9:00:51 PM**
Earliest Snapshot Time **Oct 30, 2021 10:29:07 PM**

Fig. 2.11. Edit Automatic Workload Repository.

2.4. Using advisors.

The use of advisors in monitoring is characterized by:

- for a certain specific object for analysis - provide information and recommended actions;
- are called implicitly by the Oracle Database server or explicitly by the DBA.

Wizards are procedures that a user can call (or that the Oracle Database server can call internally) that specify a specific object for analysis. The advisor can report on various aspects of the object and describe a recommended action for any condition that requires user intervention. The advisor may report that the condition can be corrected through an automated task that is provided. It is common for ADDM, or a given alert, to recommend running a specific wizard to analyze the problem in more detail.

2.4.1. Advisors available.

The following performance advisors are built into Enterprise Manager:

- **ADDM: Automatic Database Diagnostic Monitor - General Database Advisor.** Its job is to perform a top-down analysis of the system, identify problems and their potential causes, and make recommendations for troubleshooting. The overall goal is to reduce all "bottlenecks" in the system and therefore improve performance. Potentially can connect with other advisors.

- **SQL Tuning Advisor:** Analyzes SQL statements and makes recommendations to improve performance.

- **SQL Access Advisor:** Tunes a schema to a given SQL workload and provides recommendations such as creating indexes and real views for a given workload.

- **Memory Advisors:** Provides graphical analyzes of total memory target settings, System Global Area (SGA) and Program Global Area (PGA) target settings, and SGA component size settings. These analyzes are used to tune database performance and plan likely actions. Different memory wizards are available depending on the memory management mode:

- automatic memory management is enabled: only the memory advisor is available, providing advice on the overall memory target for the instance.

- automatic shared memory management is enabled: SGA Advisor and PGA Advisor are available.

- manual management of shared memory is enabled: shared pool wizard, buffer cache wizard and PGA wizard are available.

- **Segment Advisor:** Provides advice on whether an object is a good candidate for a segment shrink operation based on the level of space fragmentation in the object. The advisor also reports the growth trend of the segments in historical aspect. This information can be used for capacity planning and to reach an informed decision about which segments to shrink.

- **Undo Advisor (Undo Advisor):** Supports the sizing of the undo space of the tablespace by taking into account system activity statistics, the longest running query, and the low threshold value for retention of an action specified in the initialization parameter UNDO_RETENTION.

2.4.2. Manually calling ADDM.

ADDM runs every 60 minutes by default. Performance findings from the last snapshot are listed on the database home page.

ADDM Can be called manually. This may need to be done for the following reasons:

- If an alert-related action is recommended.
- Must start in the middle of a snapshot period.
- Must run in multiple instant states.

ADDM can be invoked manually by following these steps:

1. Select Advisor Central in the Related Links area of the database home page.
2. ADDM is selected.

The Run ADDM page appears. Increased session activity is shown as spikes in the graph.

3. Select "Run ADDM to analyze current performance" to create a new AWR snapshot and run ADDM on the new and previous snapshots. "Run ADDM to analyze past state" can be selected and start and end time of the period can be specified (Fig. 2.12.).

4. Confirm with OK.

A confirmation page appears, then a new snapshot is taken.

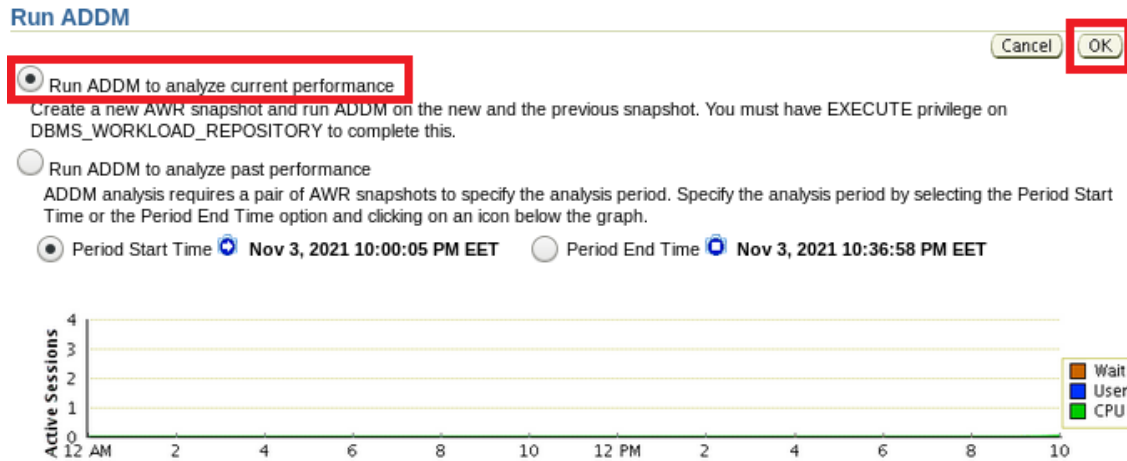


Fig. 2.12. Run ADDM to analyze current performance.

2.4.3. Usage of the SQL Tuning Advisor.

SQL Tuning Advisor can be used to analyze SQL statements and get performance recommendations. Typically, this wizard is launched as a result of an ADDM performance finding.

Additionally, SQL Tuning Advisor can be run when the top SQL statements consuming the most CPU time, I/O operations, and memory need to be analyzed.

SQL Tuning Advisor can be started in the following cases:

- Highest activity: Used to include SQL statements that may have caused recent performance problems. The wizard evaluates the most intensive SQL statements executed in the last hour.
- Historical SQL: Used to proactively tune SQL statements. The wizard evaluates a set of SQL statements during each 24-hour window.
- SQL Tuning Set (STS): The wizard evaluates a set of SQL statements that are provided.

An STS can be created from SQL statements captured from AWR snapshots or from any SQL workload.

2.4.4. Automatic SQL Tuning Advisor.

The SQL Auto Setup Advisor runs automatically when using a system maintenance window as a maintenance task. During each automatic run, the advisor selects high-load SQL queries and generates tuning recommendations for those queries. DML statements are not considered by the SQL Autotune Advisor.

One type of recommendation that the SQL Auto Setup Advisor generates is to create or modify a SQL profile. This type of recommendation can be applied automatically. Other types of recommendations (such as creating new indexes, refreshing optimizer statistics, or restructuring SQL) can only be applied manually.

Configure SQL Auto Setup Advisor.

When configuring the SQL Auto Setup Advisor, the following tasks can be performed:

- Enable automatic enforcement of recommendations for SQL profiles. Automatic deployment is disabled by default.
- Selects the maintenance windows in which the advisor runs. By default, the SQL Auto Setup advisor runs in all support windows.
- Change the start time and duration of existing maintenance windows or create new maintenance windows.

The SQL Auto Setup advisor is configured as follows:

1. The server properties page is accessed. Then Automated Maintenance Tasks in the Oracle Scheduler section.
2. Configure (Configure) page Automated maintenance tasks (Fig. 2.13.).
3. Here to disable the SQL auto-tuning advisor, select 'Disabled' for SQL auto-tuning in the Task Settings section.
4. (Optional) To prevent the SQL Auto Setup advisor from starting in a specific maintenance window, uncheck the appropriate box in the Maintenance Window Group Assignment section.

Database Instance: orcl > Automated Maintenance Tasks > Logged in As SYS

Show SQL Revert Apply


Automated Maintenance Tasks Configuration

Global Status Enabled Disabled

Task Settings

Optimizer Statistics Gathering Enabled Disabled Configure

Segment Advisor Enabled Disabled

Automatic SQL Tuning Enabled Disabled Configure 

Maintenance Window Group Assignment

Edit Window Group

Window	Optimizer Statistics Gathering	Segment Advisor	Automatic SQL Tuning
	Select All Select None	Select All Select None	Select All Select None
WEDNESDAY WINDOW	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
THURSDAY WINDOW	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>


Fig. 2.13. Configure automated maintenance tasks.

5. To enable the automatic application of recommendations for SQL profiles, select Configure in the SQL automatic setup options (Fig. 2.14) in the Task Settings section.

Automatic SQL Tuning Settings

Show SQL Revert Apply

Maximum Time Spent Per SQL During Tuning (sec)

 Automatic Implementation of SQL Profiles Yes No

Maximum SQL Profiles Implemented Per Execution

Maximum SQL Profiles Implemented (Overall)

Show SQL Revert Apply

Fig. 2.14. Automatic application of recommendations for SQL profiles.

6. Select Yes in the "Automatically deploy SQL profiles" field.
7. Apply. A confirmation message is received.

2.4.5. View SQL Autotune results.

To review the results of the SQL autotuning, perform the following sequence of actions:

1. The Server page opens. In the Oracle Scheduler / Automated Maintenance Tasks section.

The Automated Maintenance Tasks page appears.

2. Automatic SQL setting is selected. The SQL Auto-Tuning Results Summary page displays graphical summaries of the activities and findings of the SQL Auto-Tuning Advisor (Fig. 2.15.).

3. Select View Report to view recommendations. The "Automatic SQL Turing Result Details" page shows the SQL statements for which recommendations were made during the specified period.

Automatic SQL Tuning Result Summary

The Automatic SQL Tuning runs during system maintenance windows as an automated maintenance task, searching for ways to improve the execution plans of high-load SQL statements.

Task Status

Automatic SQL Tuning (SYS_AUTO_SQL_TUNING_TASK) is currently **Enabled** [Configure](#)
Automatic Implementation of SQL Profiles is currently **Disabled** [Configure](#)
Key SQL Profiles **3** [Implement All](#)

Summary Time Period

Choose a time period to focus the graphs and statistics below on a specific range of tuning results. Drill down to view focused results or see the results for all SQLs by clicking the "View Report" button.

Time Period [Go](#) [View Report](#)

Begin Date **Oct 12, 2021 11:00:10 PM (UTC+02:00)** End Date **Nov 3, 2021 11:20:45 PM (UTC+02:00)**

Overall Task Statistics

Executions **6** Candidate SQL **144** Distinct SQL Examined **46**

SQL Examined Status



Breakdown by Finding Type




Fig. 2.15. Summary of SQL Autotune results

4. Select SQL Text and View Recommendations. The SQL ID Recommendations page details each recommendation for the text. On this page, one can select a recommendation and then click on Apply to apply it.

2.4.6. Usage of the SQL Access Advisor.

SQL Access Advisor is used to tune the user schema and improve the performance of a given query. This advisor requires the user to identify a SQL workload, which is a representative set of SQL statements that access the schema. A custom workload can be selected from a variety of sources, including current and recent SQL activity, a SQL repository, or a user-defined workload (for example, from a development environment).

SQL Access Advisor can make recommendations such as creating indexes or materialized views to improve query performance for the given workload.