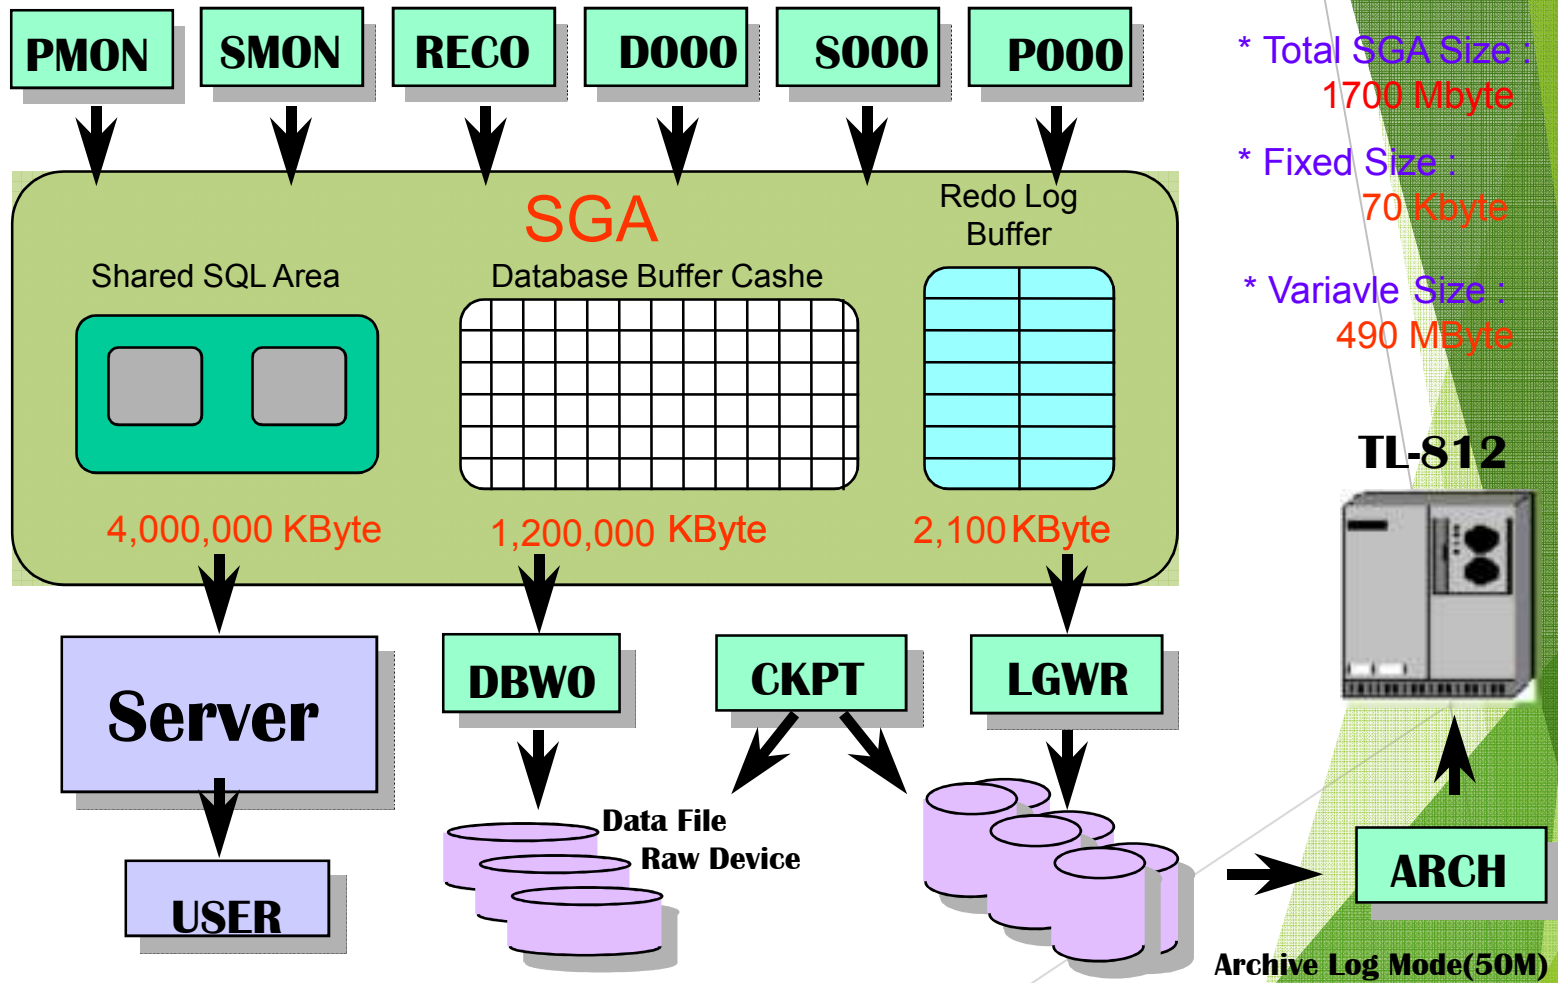# Agenda

- **Introduction on Backup & Recovery**
- **Key Data Structures of Backup and Recovery**
- **Types of Backups**
- **Various Backup methods/Tools**
- **Recovery Strategy**
- **Types of Recovery**
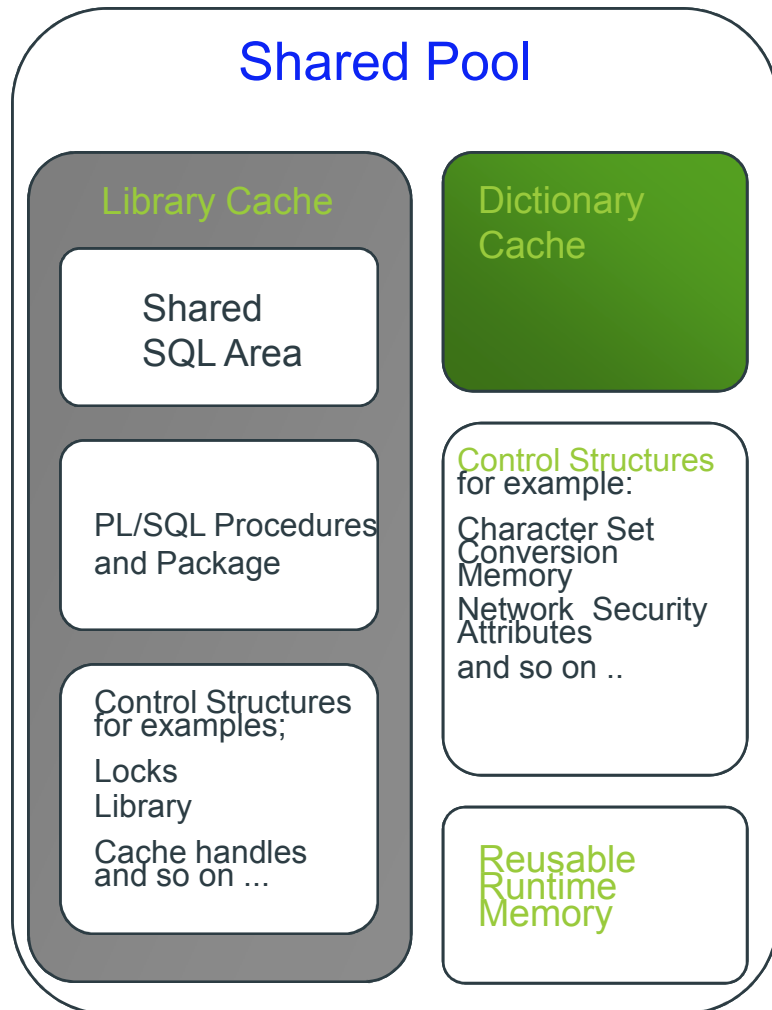- **Various Recovery Methods and Techniques**

Introduction

▶ A backup is a snapshot of a datafile, tablespace, or database at a certain time.

▶ This copy can include important parts of a database such as the control file, redo logs and datafiles

▶ A backup protects data from application error and acts as a safeguard against unexpected data loss, by providing a way to restore original data.
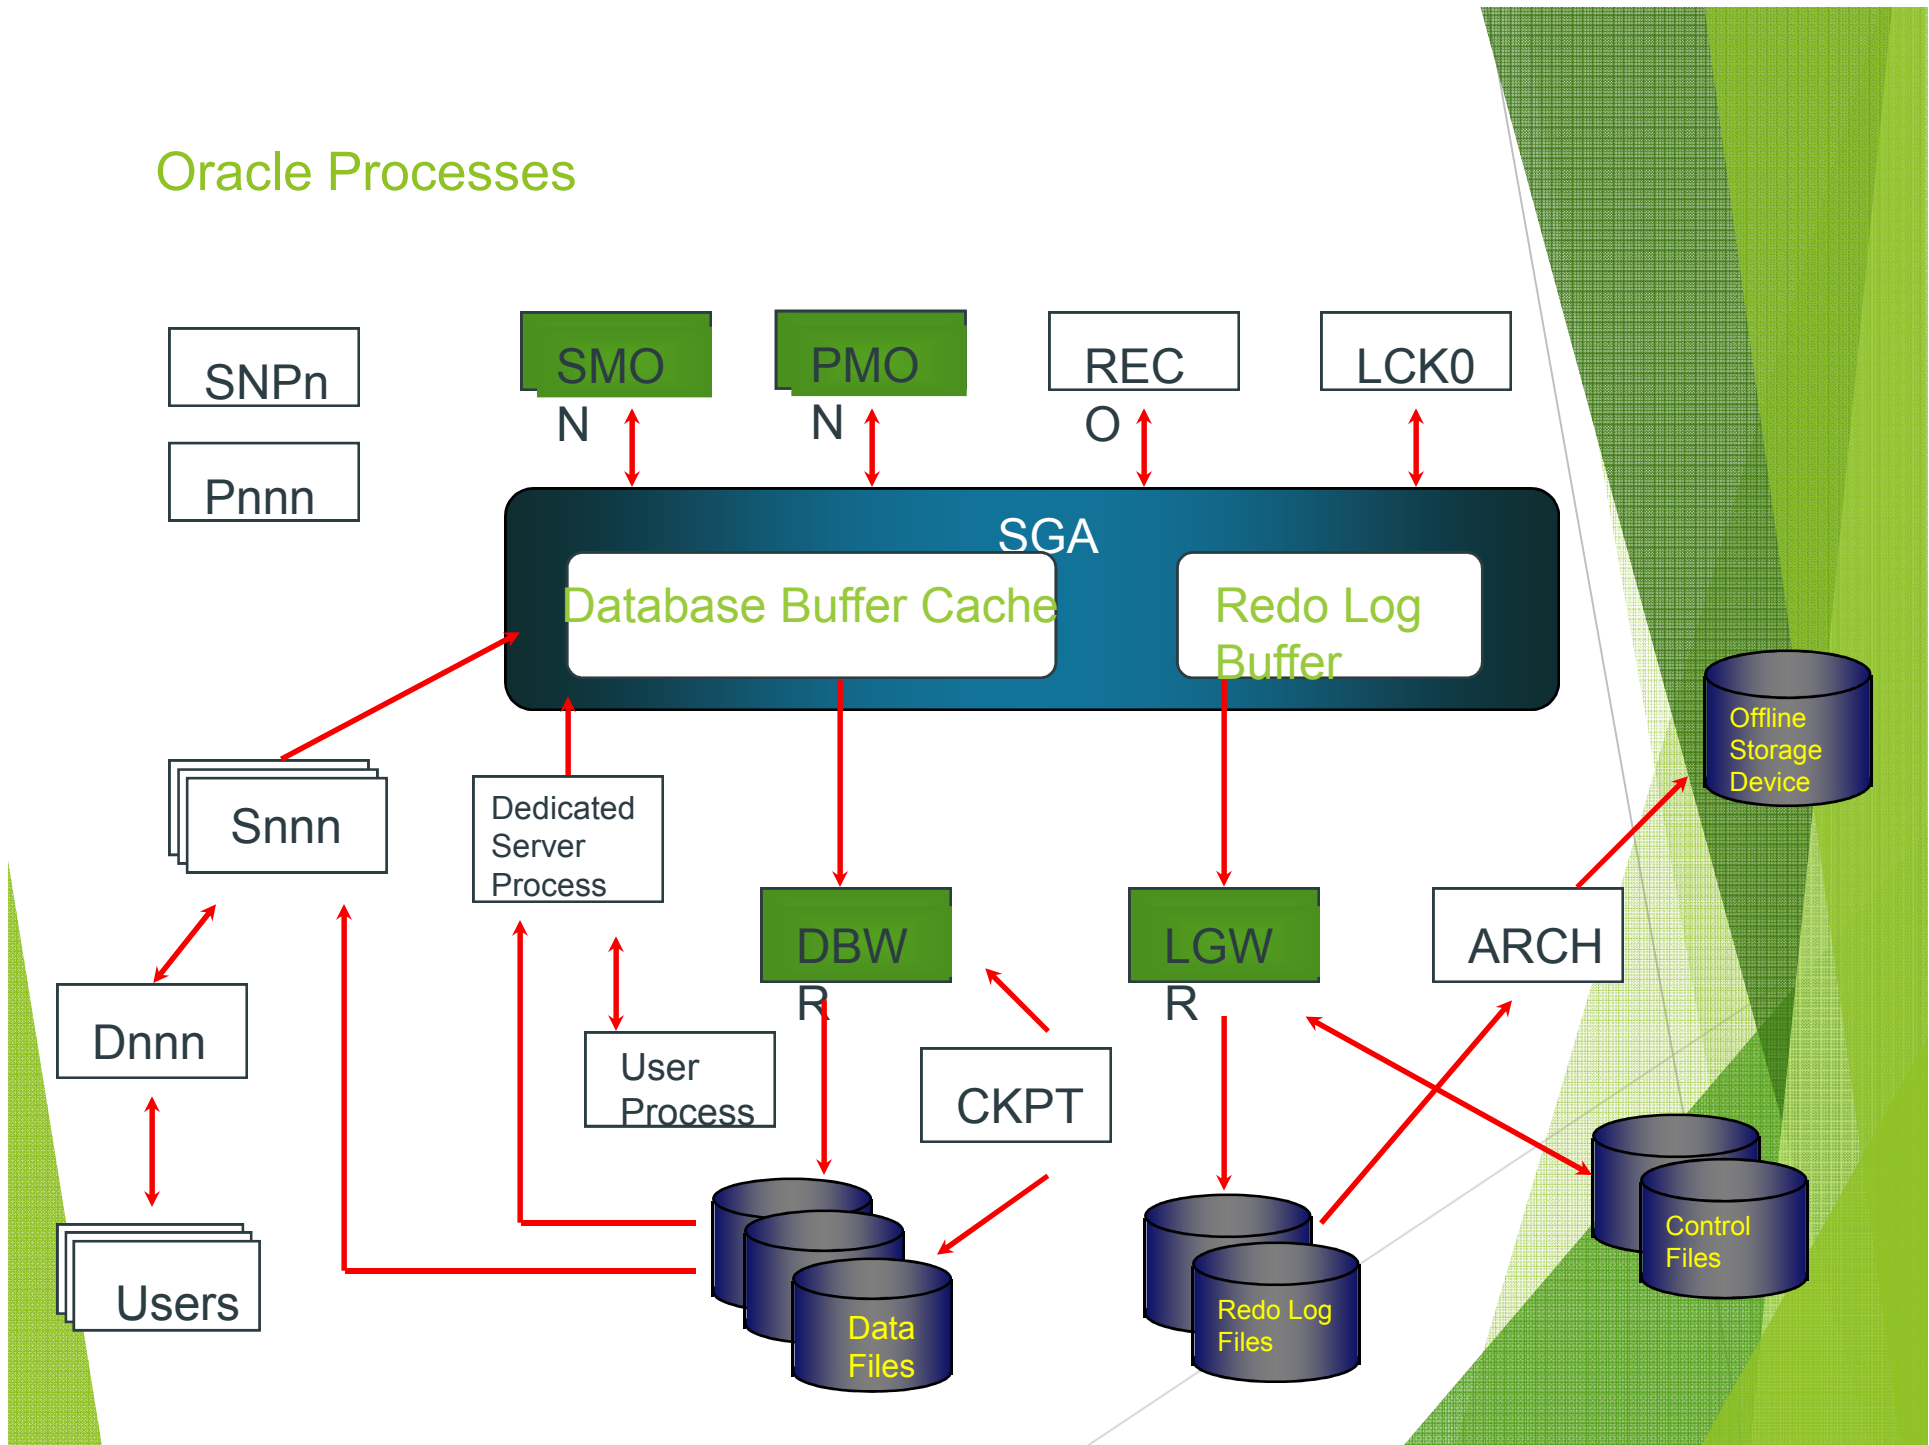
10/21/20
16

2

Overview of Oracle Architecture

PMON	SMON	RECO	D000	S000	P000

SGA

Redo Log Buffer

* Total SGA Size : 1700 Mbyte

* Fixed Size : 70 Kbyte

* Variavle Size : 490 MByte

Shared SQL Area

Database Buffer Cashe

4,000,000 KByte

1,200,000 KByte

2,100 KByte

TL-812

Server

DBW0

CKPT

LGWR

USER

Data File
Raw Device

ARCH

Archive Log Mode(50M)

Memory Structure : Shared Pool

Shared Pool

Library Cache

Shared SQL Area

PL/SQL Procedures and Package

Control Structures
for examples;
Locks
Library
Cache handles
and so on ...

Dictionary Cache

Control Structures
for example:

Character Set
Conversion
Memory

Network Security
Attributes
and so on ..

Reusable Runtime Memory

- Shared Pool Contents
  - Text of the SQL or PL/SQL statement
  - Parsed form of the SQL or PL/SQL statement
  - Execution plan for the SQL or PL/SQL statements
  - Data dictionary cache containing rows of data dictionary information

- Library Cache
  - shared SQL area
  - private SQL area
  - PL/SQL procedures and package
  - control structures : lock and library cache handles

- Dictionary Cache
  - names of all tables and views in the database
  - names and datatypes of columns in database tables
  - privileges of all Oracle users

- SHARED_POOL_SIZE

Oracle Processes

SNPn

Pnnn

SMO
N

PMO
N

REC
O

LCK0

SGA

Database Buffer Cache

Redo Log Buffer

Snnn

Dedicated Server Process

Dnnn

User Process

Users

DBW
R

CKPT

LGW
R

ARCH

Offline Storage Device

Data Files

Redo Log Files

Control Files

Oracle Database

▶ Physical Structure

   ▶ Control Files (binary files)

   ▶ Data Files

   ▶ Redo log files

Data Files

Control Files

Redo Log Files

Oracle Database
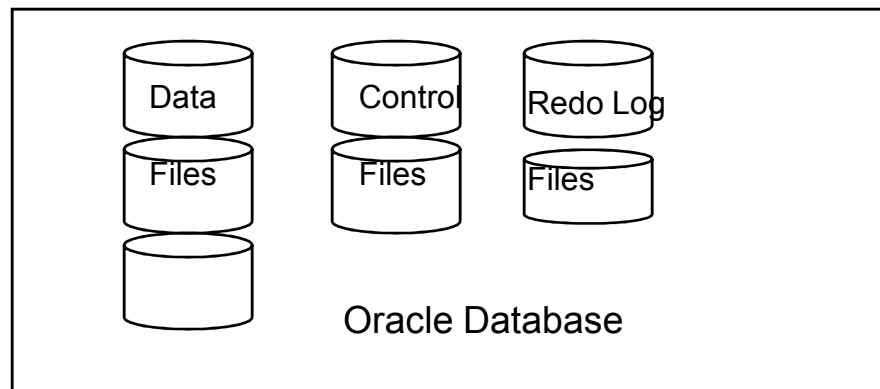
Key Data structures for backup and Recovery

- **Datafiles (*.dbf)**

  The datafiles contain all the database data. The data of logical database structures, such as tables and indexes, is physically stored in the datafiles allocated for a database.

- **Control Files (*.ctl)**

  Every Oracle database has a control file. A control file contains entries that specify the physical structure of the database such as Database name and the Names and locations of datafiles and redo log files.

- **Redo Log Files (*.log)**

  The primary function of the redo log is to record all changes made to data. If a failure prevents modified data from being permanently written to the datafiles, then the changes can be obtained from the redo log, so work is never lost.

- **Archive Log Files (*.log)**

  Oracle automatically archives log files when the database is in ARCHIVELOG mode. This prevents oracle from overwriting the redo log files before they have been safely archived to another location.

- **Parameter Files (initSID.ora)**

  Parameter files contain a list of configuration parameters for that instance and database.

  Automatic Managed undo

  Every Oracle database must have a method of maintaining information that is used to roll back, or undo, changes to the database. Such information consists of records of the actions of transactions, primarily before they are committed

ELEMENTS OF A BACKUP AND RECOVERY STRATEGY

Although backup and recovery operations can be intricate and vary from one business to another, the basic principles follow these four simple steps:

▶ Multiplex the online redo logs

▶ Run the database in ARCHIVELOG mode and archive redo logs to multiple locations

▶ Maintain multiple concurrent backups of the control file

▶ Take frequent backups of physical datafiles and store them in a safe place, making multiple copies if possible

Types of Backups

Backups are divided into physical backups and logical backups

▶ Physical backups are copies of physical database files

▶ logical backups contain data that is exported using SQL commands and stored in a binary file . They are used to supplement physical backups.

Whole Database and Partial Database Backups

▶ Whole Database Backups

  Online Database backup

  Offline Database backup

▶ Table space Backups

▶ Datafile Backups

▶ Control File Backups

Hot Backup (aka Inconsistent Backups)

▶ If the database must be up and running 24 hours a day, seven days a week, then you have no choice but to perform inconsistent backups of the whole database. A backup of online data files is called an online backup. This requires that you run your database in ARCHIVELOG mode.

▶ Advantage : Database remains open during backup

▶ Disadvantage : Large databases may have performance impact during backup, recovery takes longer and is *slightly* more complex

▶ is a backup in which all read-write datafiles and control files have not been checkpointed with respect to the same SCN.

What is offline database Backup

Cold Backup (aka Consistent Backups)

▶ The only way to make a consistent whole database backup is to shut down the database with the NORMAL, IMMEDIATE, or TRANSACTIONAL options and make the backup while the database is closed.

▶ Advantage : No recovery is required after datafiles are restored – quicker restore

▶ Disadvantage : No access to database during backup time (depends on size/system speed)

▶ In this backup, all datafiles and control files are consistent to the same point in time - consistent with respect to the same SCN,

Other Database Backups

▶ What is Tablespace Backup

  A tablespace backup is a subset of the database. Tablespace backups are only valid if the database is operating in ARCHIVELOG mode. The only time a tablespace backup is valid for a database running in NOARCHIVELOG mode is when that tablespace is read-only or offline-normal.

▶ What is Datafiles Backup

 A datafile backup is a backup of a single datafile. Datafile backups, which are not as common as tablespace backups and are only valid if the database is run in ARCHIVELOG mode. The only time a datafile backup is valid for a database running in NOARCHIVELOG mode is if that datafile is the only file in a tablespace. For example, the backup is a tablespace backup, but the tablespace only contains one file and is read-only or offline-normal. .

▶ What is Controlfile Backup

  A control file backup is a backup of a database's control file. If a database is open, the user can create a valid backup by issuing the following SQL statement: ALTER DATABASE BACKUP CONTROLFILE to 'location'; or use Recovery Manager (RMAN).

▶ What is Archived redo log Backup

Archived redo logs are the key to successful media recovery. Depending on the disk space available and the number of transactions executed on the database, you want to keep as many days of archive logs on disk and you want to back them up regularly to ensure a more complete recovery

Logical backups

▶ Logical backup (Export)

  ▶ Logical backups are exports of schema objects, like tables and stored procedures, into a binary file. Oracle utilities are used to move Oracle schema objects in and out of Oracle.

  ▶ Not recommended for backup of a whole database, but useful for backing up individual objects or schemas or moving data into another database

  ▶ It is platform independent.

Example

exp scott/tiger file=emp.dmp log=emp.log tables=emp rows=yes indexes=no exp scott/tiger file=emp.dmp tables=(emp,dept)

imp scott/tiger file=emp.dmp full=yes imp scott/tiger file=emp.dmp fromuser=scott touser=scott tables=dept

Backup Methods :

Oracle provides users a choice of several basic methods for making backups. The methods include:

▶ Recovery Manager (RMAN) - A component that establishes a connection with a server process and automates the movement of data for backup and recovery operations.

▶ Oracle Enterprise Manager - A GUI interface that invokes Recovery Manager.

▶ Oracle Data Pump - The utility makes logical backups by writing data from an Oracle database to operating system files in a proprietary format. This data can later be imported into a database.

▶ User Managed - The database is backed up manually by executing commands specific to the user's operating system.

RMAN backup

▶ Saves storage space, because it only backs up blocks that have ever held data.

▶ Eliminates performance overhead during backups, because it doesn't place tablespaces in backup mode.

▶ Helps maintain database integrity by detecting block corruption during backups.

▶ Decreases backup time by allowing parallelized and incremental backups.

▶ Maintains history of old backups and status of running backups.

▶ Has built in reporting capabilities.

▶ Supports incremental backup strategies

▶ Can seamlessly work with third party media managers

▶ Allows transition to 10^g Automatic Storage Management which requires RMAN

▶ Decreases time for specific types of recoveries by allowing block and tablespace point in time options.

Examples

$Connect to RMAN

CONFIGURE DEVICE TYPE sbt PARALLELISM 1;

CONFIGURE CHANNEL DEVICE TYPE sbt

backup database;

backup archivelog from time 'sysdate-2' skip inaccessible;

backup current controlfile;

list backup summary;

list backup;

SQL> ALTER SESSION SET nls_date_format = 'DD-MON-YY HH24:MI';

SQL>SELECT start_time, completion_time, status, keep_until, backup_type

      FROM rcat_dbtst.rc_backup_set@rmand.corp.company.com

  ORDER BY start_time

/

User managed backup

▶ You must use operating system utilities to make user-managed backups.

▶ The available commands are operating system specific. For example, on a UNIX system

you can back up a datafile using cp as follows:

cp /oracle/dbs/df1.f /backup/df1.f

▶ One major difference between user-managed backups and RMAN backups is that in the former there is no automatic metadata record of the backup. In other words, you must manually keep records of what you back up and where you back it up.

▶ For hotbackup,database has to be in Hotbackup mode before taking the backup using cp command.

▶ Database generates extra redo while hotbackup is active .The reason for extra redo generation is fractured block -- the head and tail are from two points in time.

Difference in Manual hotbackup and RMAN hotbackup

▶ We need to put database in hotbackup mode in manual backup while RMAN does not require database to be in hotbackup mode.

▶ Extra redo is generated while taking the hotbackup using manual backup while no overhead is there with RMAN backup

▶ Datafiles header are not updated with latest checkpoint while taking hotbackup using manual mode while with RMAN,it keeps on updating

Recovery Strategy

▶ Basic recovery involves two parts:

▶ restoring a physical backup and

▶ then updating it with the changes made to the database since the last backup.

▶ The most important aspect of recovery is making sure all data files are consistent with respect to the same point in time

When preparing a recovery strategy, it is critical to understand the answers to these questions:

▶ How does recovery work?

▶ What are the types of recovery?

▶ Which recovery method should be used?

Types of Recovery

There are 3 types of Recovery :

▶ Instance Recovery

▶ Crash Recovery

▶ Media Recovery

Oracle performs the first two types of recovery at instance startup. Only media recovery requires the user to issue commands. Instance recovery happens only in case of RAC database where surviving instance recovers the failed instance by using the redo logs to recover the committed data in the database buffers. Oracle also undo the uncommitted transactions using undo segments.

Media Recovery

Complete Recovery

In this recovery, database can be recovered to the exact point of failure. It is called complete because Oracle applies all the redo changes to the backup

Incomplete Recovery

if a database can not be recovered to the most current time as we do not have sufficient archived redo logs, then database can be recovered to a point of time until we have the archived redo logs.

Incomplete Recovery Options :

▶ In case of incomplete recovery, the user can perform Tablespace point-in-time recovery (TSPITR), which enables users to recover one or more tablespaces to a point-in-time that is different from the rest of the database.

▶ Time-based recovery, also called point-in-time recovery (PITR), which recovers the data up to a specified point in time.

▶ Cancel-based recovery, which recovers until the CANCEL command is issued.

▶ Change-based recovery or log sequence recovery. If O/S commands

are used, change-based recovery recovers up to a specified SCN in the redo record.

▶ Flashback from human error

Recovering from Human Error

► The Oracle Database 10g architecture uses flashback to recover from human errors.

► Flashback Technology provides a set of new features to view and rewind data back and forth in time.

► The Flashback features offer the capability to query past versions of schema objects, query historical data, perform change analysis or perform self-service repair to recover from logical corruptions.

Recovery using RMAN

▶ The basic RMAN commands are RESTORE and RECOVER.

▶ RMAN can be used to restore datafiles from backup sets or image copes, either to their current location or to a new location. If any archived redo logs are required to complete the recovery operation, RMAN automatically restores and applies them.

▶ In a recovery catalog, RMAN keeps a record containing all the essential information concerning every backup ever taken. If a recovery catalog is not used, RMAN uses the control file for necessary information.

▶ The RMAN RECOVER command can be used to perform complete media recovery and apply incremental backups, and to perform incomplete media recovery

Recovery Example

DATAFILE DROPPED

SQL> SELECT COUNT(*) FROM plan_table;   **-- 10 rows**

tech:/oracle/tools/rman> rm -i /u01/oracle/test/dbtst_tools1.dbf   **-- remove file**

SQL> SELECT COUNT(*) FROM plan_table;

SELECT COUNT(*) FROM plan_table

                    *****

ERROR at line 1:

ORA-01116: error in opening database file 10

ORA-01110: data file 10: `/u01/oracle/test/dbtst_tools1.dbf **'**

ORA-27041: unable to open file

$Connect to RMAN

shutdown immediate;

shutdown abort;

startup mount;

run {

allocate channel t1 type 'SBT_TAPE`;

restore datafile 10;

recover datafile 10;

alter database open;

}

SQL> SELECT COUNT(*) FROM plan_table;   **-- 10 rows**

Recovery Using sqlplus

▶ Administrators can use the SQL*Plus utility at the command line to restore and perform media recovery on your files. Users can recover a   database, tablespace, or datafile. Before performing recovery, users need to:

▶ Determine which files to recover. Often the table
 V$RECOVER_FILE can be used.

▶ Restore backups of files permanently damaged by media failure.

▶ If the user does not have a backup, recovery can still be Performed if the user has the necessary redo log files and the control file contains the name of the damaged file.

▶ If a file cannot be restored to its original location, then the user must relocate the restored file and inform the control file of the \ new   location.

▶ Restore necessary archived redo log files. .

Conclusion :

▶ Backup and recovery of your Oracle database is important to protecting data from corruptions, hardware failures, and data failures. While Oracle provides many features to protect your data, nothing can replace a backup.