Managing Redo Log Files

Redo logs consist of two or more pre allocated files that store all changes made to the database. Every instance of an Oracle database has an associated online redo log to protect the database in case of an instance failure.

Main points to consider before creating redo log files?

- Members of the same group should be stored in separate disks so that no single disk failure can cause LGWR and database instances to fail.
- Set the archive destination to separate disk other than redo log members to avoid contention between LGWR and Arch.
- With mirrored groups of online redo logs , all members of the same group must be the same size.

What are the parameters related to Redo log files?
Parameters related to redo log files are

- MAXLOGFILES
- MAXLOGMEMEBERS

MAXLOGFILES and ***MAXLOGMEMEBERS*** parameters are defined while creation of database. You can increase these parameters by recreating the control file.

Activity

How do you create an online Redo log group?

Alter database add logfile group <group Number> ('<DISK>:\Directory\<LOG_FILE_NAME>.log','
('<DISK>:\Directory\<LOG_FILE_NAME>.log') size 500K;

How to check the status of added redo log group?

Select * from v$log;

Interpretation:
Here you will observe that status is UNUSED means that this redo log file is not being used by oracle as yet. ARC is the archived column in v$log , it is by default YES when you create a redo log file. It will returns to NO if the system is not in archive log mode and this file is used by oracle. Sequence# 0 also indicate that it is not being used as yet.

How to create online redo log member ?

alter database add logfile member '<DISK>:\Directory\<LOG_FILE_NAME>.log',<DISK>:\Directory\<LOG_FILE_NAME>.log'' to group <GROUP NUMBER>;

How to rename and relocate online redo log members ?

Important: Take the backup before renaming and relocating.
Step:1 Shutdown the database .
Step:2 Startup the database as startup mount.
Step:3 Copy the desired redo log files to new location . You can change the name of redo log file in the new location.
Step:4 Alter database rename file '<DISK>:\Directory\<LOG_FILE_NAME>.log' to '<new path><DISK>:\Directory\<LOG_FILE_NAME>.log',
Step:5 Alter database open;
Step: 6 Shutdown the database normal and take the backup.

How to drop an online redo log group?

Important:

- You must have at- least two online groups.
- You can not drop a active online redo log group. If it active switch it by alter system switch logfile before dropping.
- Also make sure that online redo log group is archived (if archiving is enabled).

Syntax:
If you want to drop log group:

Alter database drop logfile group <GROUP_NUMBER>;

If you want to drop a logfile member:

Alter database drop logfile member ' <DISK>:\Directory\<LOG_FILE_NAME>.log';

How to View Online Redo Log Information?

SELECT * FROM V$LOG;

GROUP# THREAD# SEQ BYTES MEMBERS ARC STATUS FIRST_CHANGE# FIRST_TIM

------ ------- ----- ------- ------- --- --------- ------------- ---------

1 1 10605 1048576 1 YES ACTIVE 11515628 16-APR-00

2 1 10606 1048576 1 NO CURRENT 11517595 16-APR-00

3 1 10603 1048576 1 YES INACTIVE 11511666 16-APR-00

4 1 10604 1048576 1 YES INACTIVE 11513647 16-APR-00

SELECT * FROM V$LOGFILE;

GROUP# STATUS MEMBER

------ ------- ---------------------------------

1 D:\ORANT\ORADATA\IDDB2\REDO04.LOG

2 D:\ORANT\ORADATA\IDDB2\REDO03.LOG

3 D:\ORANT\ORADATA\IDDB2\REDO02.LOG

4 D:\ORANT\ORADATA\IDDB2\REDO01.LOG

If STATUS is blank for a member, then the file is in use.

FAQ Related to Redo Log Files

Explain the purpose of Oracle Redo Log files in a database, and discuss their significance for database recovery and transaction consistency.

Oracle Redo Log files are critical components of the Oracle database architecture, designed to ensure data integrity, transaction durability, and database recoverability.

Redo Log files record all changes made to the database, capturing a chronological sequence of data modifications (inserts, updates, and deletes). These changes are stored in a redo log format, providing a detailed log of all committed transactions.

Redo Log files play a crucial role in both instance recovery and media recovery. In the event of an instance or media failure, the redo log allows the database to roll forward or roll back transactions to bring the database to a consistent state.

The redo log includes information about committed transactions (commit records) and aborted transactions (rollback records). This information is vital for maintaining the ACID properties of transactions, ensuring that committed changes are durable and consistent.

Redo Log files are typically configured in groups, and each group consists of multiple members (individual redo log files). This configuration allows for cyclical use of the log files and helps in minimizing the risk of data loss.

In an Oracle database, the Redo Log files can be configured to operate in either ARCHIVELOG mode or NOARCHIVELOG mode. ARCHIVELOG mode enables the archiving of filled redo log files, providing a point-in-time recovery option.

How does the ARCHIVELOG mode differ from NOARCHIVELOG mode, and when would you choose one over the other?

In archive log mode, the database enables the archiving of filled Redo log Files. It Allows point-in-time recovery, providing the ability to recover the database to a specific point in time. It is Recommended for production databases where data loss is unacceptable.

In noarchiveog mode, the database does not archive Redo Log files; they are overwritten when filled. So we can not perform the recovery. It is Suitable for non-production environments or scenarios where point-in-time recovery is not a critical requirement.

Can you explain the concept of Log Switch and its role in managing Redo Log files?

A log switch is the process of Oracle transitioning from one online Redo Log file to another within a group. As database changes occur, they are recorded in the current Redo Log file. Once the file is full or a log switch is triggered, Oracle starts writing to the next available Redo Log file in the same group.

What happens if a Redo Log file group is corrupted or lost? How can recovery be performed?

Database functionality may be compromised, and recovery is necessary.
Recovery Steps
Determine the corrupted or lost Redo Log file group.
Restore the affected group from a valid backup.
If needed, recreate the control file from a backup.
Open the database with the RESETLOGS option.
Use recovered Redo Log files to roll forward or roll back.
Test and verify the consistency of the recovered database.

How can you add or drop Redo Log file groups in an Oracle database?

To add a Redo Log file group, use the `ALTER DATABASE` statement with the `ADD LOGFILE` clause.
To drop a Redo Log file group, ensure that the group is not active.
Use the `ALTER DATABASE` statement with the `DROP LOGFILE GROUP` clause.
After dropping, consider adding new Redo Log file groups to maintain availability.

Discuss the significance of sizing and multiplexing Redo Log files for performance and reliability.

Adequate sizing ensures that Redo Log files are large enough to handle transactional workloads without frequent log switches. This minimizes performance overhead associated with frequent file switches.
Multiplexing involves creating redundant copies of Redo Log files on different physical devices or locations. This guards against data loss due to a failure in one copy, enhancing overall system reliability.