

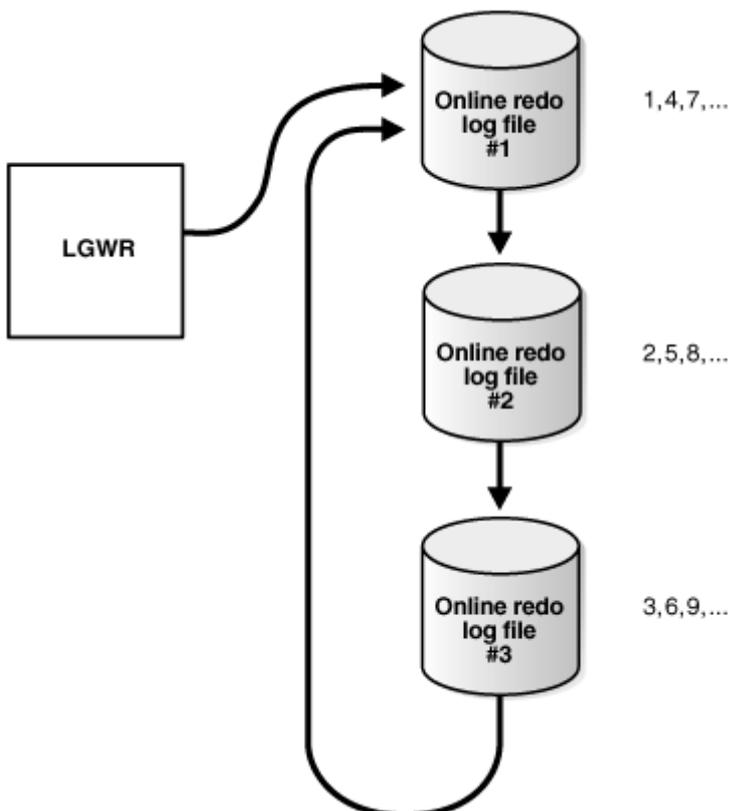
Oracle Notes Part-2

Redo: In the Oracle RDBMS environment, **redo logs** comprise files in a proprietary format which log a history of all changes made to the database. Each redo log file consists of redo records. A redo record, also called a redo entry, holds a group of change vectors, each of which describes or represents a change made to a single block in the database.

For example, if a user **UPDATEs** a salary-value in a table containing employee-related data, the DBMS generates a redo record containing change-vectors that describe changes to the data segment block for the table. And if the user then **COMMIT** the update, Oracle generates another redo record and assigns the change a "system change number" (SCN).

LGWR writes to redo log files in a circular fashion. When the current redo log file fills, LGWR begins writing to the next available redo log file. When the last available redo log file is filled, LGWR returns to the first redo log file and writes to it, starting the cycle again. The numbers next to each line indicate the sequence in which LGWR writes to each redo log file.

Reuse of Redo Log Files by LGWR:



Oracle Database uses only one redo log files at a time to store redo records written from the redo log buffer. The redo log file that LGWR is actively writing to is called the current redo log file. Redo log files that are required for instance recovery are called active redo log files. Redo log files that are no longer required for instance recovery are called inactive redo log files.

A log switch is the point at which the database stops writing to one redo log file and begins writing to another. Normally, a log switch occurs when the current redo log file is completely filled and writing must continue to the next redo log file. However, you can configure log switches to occur at regular intervals, regardless of whether the current redo log file is completely filled. You can also force log switches manually.

Oracle Database assigns each redo log file a new log sequence number every time a log switch occurs and LGWR begins writing to it. When the database archives redo log files, the archived log retains its log sequence number. A redo log file that is cycled back for use is given the next available log sequence number.

UNDO: Oracle Database creates and manages 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. These records are collectively referred to as undo.

Undo records are used to:

- Roll back transactions when a ROLLBACK statement is issued
- Recover the database
- Provide read consistency
- Analyze data as of an earlier point in time by using Oracle Flashback Query
- Recover from logical corruptions using Oracle Flashback features.

	Undo	Redo
Record of	How to undo a change	How to reproduce a change
Used for	Rollback, read-consistency	Rolling forward database changes
Stored in	Undo segments	Redo log files
Protects against	Inconsistent reads in multiuser systems	Data loss