

SHOP. CONNECT. ENJOY.



amazon

Privacy



Craig S. Mullins


[Return to Home Page](#)

September 2005



The DBA Corner

by Craig S. Mullins



The Log is the Database

The database log, sometimes referred to as the transaction log, is a fundamental component of a database management system. All changes to application data in the database are recorded serially in the database log. Using this information the DBMS can track which transaction made which changes to the database. Furthermore, rollback and recovery operations utilize the database log to reset the database to a particular point-in-time.

The manner in which the database log is created varies depending on the DBMS being used. Some DBMSs specify the log at the database system level, others

define a database log for each database that is created within the database system. Depending on the DBMS the database log may consist of several files. For example, Oracle uses a transaction log and rollback segments.

Some DBMSs provide a parameter to enable and disable logging. In general, avoid disabling database logging for any database or database system where the data being managed is valuable. In other words, consider turning off database logging only for test databases (and be careful even then).

During normal database application processing SQL inserts, updates, and deletes will be issued to modify data in the database. As these database modifications are made, the transaction log for the database will grow. Since each database change is logged, the DBA will need to actively monitor the size of the transaction log files. Since data is constantly changing, the log will be continuously growing.

Really, the database log is the most important component of the DBMS. Another way to think about it is that the log is the database. There is actually more information in the log because it contains a record of every change made to the database. The database itself can be thought of as just an optimized access path to the most recent version of the log.

How the Log Works

A database transaction log is a write-ahead log. This means that changes are made to the transaction log before they are actually made to the data in the database tables. When the database modification has been fully recorded on the log, recovery of the transaction is guaranteed.

Typically a system checkpoint is taken by the DBMS to guarantee that all log records and all modified database pages are written safely to disk. The frequency

of database system checkpoints can be set up by the DBA using database configuration parameters – usually checkpoint frequency is set either as a predetermined time interval or as a preset number of log records written.

Generally, the following type of information is recorded on the database log:

- the beginning and ending time of each transaction
- the actual changes made to the data and enough information to undo the modifications made during each transaction (accomplished using before and after images of the data)
- the allocation and deallocation of database pages
- the actual commit or rollback of each transaction

Using this data the DBMS can accomplish data integrity operations to ensure consistent data is maintained in the database. The transaction log is used when the DBMS is restarted, when transactions are rolled back, and to restore a database to a prior state. Let's examine each of these scenarios.

When the DBMS is restarted, each database goes through a recovery process. During restart processing, the DBMS will check to determine which transactions must be rolled forward. This occurs for transactions where it is unknown if all the modifications were actually written from the cache to disk. A checkpoint forces all modified pages to disk. Therefore, it represents the point at which the startup recovery must start to roll transactions forward. Because all pages modified before the checkpoint are guaranteed to be written accurately to disk, there is no need to roll forward anything done before the checkpoint.

When a transaction is rolled back the DBMS copies before images to the database for every modification made since the transaction began.

During a recovery scenario the DBA can use the transaction log to restore a database. First, a backup copy of the database must be restored and then subsequent transaction log backups can be restored. This causes a roll forward of the transaction log. During a roll forward the DBMS will copy after images of each modification to the database. Using the logged data the DBMS ensures that each modification is applied in the same order that it originally occurred.

You can see where the transaction log is a useful item to have around in case of database errors, transaction errors, and to ensure data integrity. And effective setup and management of your database logs is a critically important aspect of database administration.

From [Database Trends and Applications](#), July 2005.

© 2005 Craig S. Mullins, All rights reserved.

[Home.](#)