

Oracle Rdb *Journal*



Introducing the Oracle Rdb LogMiner™

**An article from the Rdb Journal
By Norm Lastovica
January 15, 2000**

Copyright © 2000 Oracle Corporation. All Rights Reserved.

Introducing the Oracle Rdb LogMiner™



Norm Lastovica
Oracle Rdb Engineering
nlastovi@us.oracle.com

January 15, 2000




Oracle Rdb after-image journal (AIJ) files contain a wealth of useful information about the activities and history of an Oracle Rdb database. AIJ files contain all of the data needed to perform database recovery. They also contain a history of changes to data and metadata in the database, including important information that is not available from any other source. Providing record data from the AIJ file contents can be accomplished at little or no performance cost and no application changes, since Oracle Rdb already maintains the needed information in the AIJ files.

Customers have asked for a mechanism that efficiently allows access to Oracle Rdb data that has been added to, modified in, or removed from the database; that puts this data in a form that can easily be used by an application or loaded into another database; and that can do this with little or no impact on performance or applications. They want high performance data extraction, and the ability to move data rapidly from one system to another—without slowing their production systems.

Oracle is pleased to announce an exciting new mechanism for Oracle Rdb: the Oracle Rdb LogMiner. This technology extracts changed data records from an Oracle Rdb after-image journal (AIJ) file for analysis or replication to another database or application. It permits user data records that are added, modified, or deleted to be returned from an Oracle Rdb after-image journal file. You don't have to make any application changes to take advantage of LogMiner capabilities. Data records are returned either to an output "stream" (i.e. file, pipe, mailbox) or to an application 'callback' (a routine in a user-supplied sharable image).

The Oracle Rdb LogMiner gives you an interface to the data record contents of Oracle Rdb after-image journal (AIJ) files. Data records that are added, updated, or deleted by committed transactions may be extracted (unloaded) from the AIJ files in a format ready to be loaded into another database or consumed directly by your own programs.

Utilizing the Oracle Rdb LogMiner feature, you can extract changed data to:

-  Create a table of changes to a source database in the same or another database.
-  Analyze the changes using an application you've written.
-  Analyze the contents of a journal file in a very convenient and easy-to-use form.

- 🔄 Replicate changes in one database to another database with different logical and physical design.
- 🔄 Perform hours of database maintenance—with only minutes of database downtime.

Because Oracle Rdb already maintains the contents of the AIJ, record data from AIJ file contents can be extracted at little or no performance cost. Simply by checking all user data records that have been added, modified, or deleted in the after-image journal backup, LogMiner extracts all change data (representing the effects of transactions modifying the database) without having to scan the database for those records that have been changed.

To do this, a command line interface translates the binary data record contents of an AIJ file into a documented output format. Changed records in the specified database table(s) for committed transactions are extracted to an output stream (such as a disk file, mailbox, pipe, or application callback) in the order that the transactions were committed.

You can write application programs to process this extracted data stream. These applications could, for example, update a reporting database based on record changes made to an OLTP database. Or the RMU/LOAD command could be utilized to create a "change table," containing the modifications made to a table.

Operation

The new RMU /UNLOAD /AFTER_JOURNAL command translates the binary record contents of an after-image journal (AIJ) file into an output file. Data records for the specified table(s) for committed transactions are extracted to an output stream (file, device, or application callback) in the order that the transactions were committed. Data records extracted from the AIJ file are those records that transactions added, modified, or deleted in base database tables. Index nodes, database metadata, segmented string (blob), views, COMPUTED BY columns, and temporary tables are not unloadable from after-image journal files.

Only the final contents of a record for each committed transaction are extracted. Multiple changes to a single record within a transaction are condensed so that only the last revision of the record appears in the output stream. It is not possible to determine which column(s) were changed in a data record directly from the after-image journal file. The record itself would have to be compared with a prior record content in order to determine which column(s) were changed.

The database used to create the after-image journal files being extracted must be available during the RMU /UNLOAD /AFTER_JOURNAL command execution. The database is used to obtain metadata information (such as table names, column counts, record version, record compression, etc.) needed to extract data records from the AIJ. The metadata reference database may be accessed either locally (on the same computer system) or remotely (over a network connection). The database is only used as a metadata reference, is read in a single read-only transaction solely to load the metadata, and is then detached.

The after-image journal file(s) are processed sequentially. All specified tables are extracted in one pass through the after-image journal. As each transaction commit record is processed from the AIJ file, all modified and deleted records for the specified table(s) are sorted to remove duplicates and then written to the output stream(s). Transactions that rolled back are discarded. Index nodes, system records, metadata,

and data records for tables not being extracted are discarded.






In the extracted output, records that were updated or added are noted as having been modified. It is not possible to distinguish between inserted and updated records in the output stream. A transaction that modifies and deletes a record generates only a delete record. A transaction that adds a new record to the database and then deletes it within the same transaction also generates only a delete record. Deleted (erased) records are returned as the final record content at the time it was deleted. A record that is modified several times before being deleted is returned as the record contents just prior to deletion.

Setting Up a Database to Allow Oracle Rdb LogMiner Extraction

Databases must be enabled for Oracle Rdb LogMiner operation. When enabled, the Oracle Rdb database software writes additional information to the after-image journal file when records are deleted from the database. This information is used during the unload operation. The new `RMU /SET LOGMINER /ENABLE` command alters a database so this additional information is automatically logged. Enabling a database for Oracle Rdb LogMiner is an off-line operation. The database and after-image journal files should be backed up afterwards.

Information Returned by Oracle Rdb LogMiner

Oracle Rdb LogMiner appends a number of data fields to the beginning of each output record. These LogMiner fields are followed by fields corresponding to the extract table's columns. This information may or may not be required by all applications and readers of the data. The information added by the Oracle Rdb LogMiner includes:

-  "M"odify or "D"elete indicator (indicates disposition of record).
-  Table name.
-  Transaction start and commit times.
-  TSN of transaction modifying the record.
-  Data record contents.

The data record itself is extracted in a format similar to that provided by the `RMU/UNLOAD/RECORD_DEFINITION=FILE` command. Fields are fixed length "binary" data (no conversion to text of integer fields, for example). `COMPUTED BY` fields are not extracted as they are not actually stored in the database.

Information That Is Not Extracted by Rdb LogMiner

Segmented string (also called "BLOB") data is not extracted. There is no link between the segments themselves and the original data record. It is possible that the segmented strings were modified though the original data record was not. This makes it very difficult to associate the changed data in the after-image journal file with the original record. Changes to temporary tables are not journaled at all, and are thus unable to be extracted. `DROP` or `TRUNCATE` table operations do not journal the actual records being erased (which is one reason these operations are relatively fast). Because the records are not journaled, they are not extracted.

Performance

As the Oracle Rdb LogMiner is currently in field test, final performance testing has not been done. However, initial testing indicates that the extraction performance should be quite good. Using a single node test system, the Oracle Rdb LogMiner software unloads one million rows from 10 tables modified by 10,000 transactions into 10 output files in about 35 seconds. In another test, an industry-standard benchmark provided the source data of about 800,000 records from eight tables and 60,000 transactions. Extracting these eight tables of change information into eight output files required about 40 seconds of elapsed time on the same system.

Examples

The following command simultaneously unloads the SALES, STOCK, SHIPPING and ORDERS tables from the AIJ backup files MFS.AIJBCK_1-JUL-1999 through MFS.AIJBCK_3-JUL-1999. Note that the input AIJ backup files are processed sequentially in the order specified.

```
$ RMU /UNLOAD /AFTER_JOURNAL MFS.RDB -
MFS.AIJBCK_1-JUL-1999, -
MFS.AIJBCK_2-JUL-1999, -
MFS.AIJBCK_3-JUL-1999 -
/TABLE = (NAME = SALES, OUTPUT = SALES.DAT) -
/TABLE = (NAME = STOCK, OUTPUT = STOCK.DAT) -
/TABLE = (NAME = SHIPPING, OUTPUT = SHIPPING.DAT) -
/TABLE = (NAME = ORDER, OUTPUT = ORDER.DAT)
```

In order to unload data based on a time range, use the /BEFORE and /SINCE qualifiers. The following example extracts changes made to the PLANETS table by transactions that committed between 1-SEP-1999 at 14:30 and 3-SEP-1999 at 16:00.

```
$ RMU /UNLOAD /AFTER_JOURNAL MFS.RDB MFS.AIJBCK -
/TABLE = (NAME = PLANETS, OUTPUT = PLANETS.DAT) -
/SINCE = "1-SEP-1999 14:30:00.00" -
/BEFORE = "3-SEP-1999 16:00:00.00"
```

To move transaction data from one database into a "change table" in another Oracle Rdb database, you can use the RMU /UNLOAD /AFTER_JOURNAL command followed by RMU /LOAD commands. A record definition (.rrd) file would need to be created for each table being loaded into the target database. The record definition files can be created using the Oracle Rdb LogMiner by specifying the RECORD_DEFINITION keyword on the TABLE qualifier.

```

$ RMU /UNLOAD /AFTER_JOURNAL OLTP.RDB MYAIJ.AIJBCK -
  /TABLE = ( NAME = MYTBL, -
            OUTPUT = MYTBL.DAT, -
            RECORD_DEFINITION=MYLOGTBL) -
  /TABLE = ( NAME = SALE, -
            OUTPUT=SALE.DAT, -
            RECORD_DEFINITION=SALELOGTBL)
$ RMU /LOAD WAREHOUSE.RDB MYLOGTBL MYTBL.DAT -
  /RECORD_DEFINITION = FILE = MYLOGTBL.RRD
$ RMU /LOAD WAREHOUSE.RDB SALELOGTBL SALE.DAT -
  /RECORD_DEFINITION = FILE = SALELOGTBL.RRD

```

Potential Use to Reduce Application Downtime

Lengthy off-line application or database maintenance operations can pose a significant difficulty in high-availability production environments. The Oracle Rdb LogMiner feature can help reduce downtime to just minutes.

If a backup copy of the database is used for maintenance operations, the application can continue running (and modifying the database) during the lengthy maintenance. Once the maintenance is complete, the application can be shut down, the production system AIJ file(s) can be backed up, and the Oracle Rdb LogMiner can be used to extract changes made to production tables since the database was backed up. These changes can then be applied (using an application or via a trigger technique as described above) to the new database. Once the new database has been brought back up to date, the application can be restarted using the new database.

The sequence of events required would be similar to the following:

1. Perform an on-line quiet-point backup of the production database.
2. Restore the backup to create a new database that will eventually become the production database.
3. Perform maintenance operations on new database (note that the production system continues to run).
4. Perform an on-line quiet-point AIJ backup of the production database.
5. Use the RMU/UNLOAD/AFTER_JOURNAL command to unload all database tables into individual output files from the AIJ backup file.
6. Using either the trigger technique or an application program, update the tables in the new database with the change data.
7. Shut down the production application and close the database.
8. Perform an off-line quiet-point AIJ backup.
9. Use the RMU/UNLOAD/AFTER_JOURNAL command to unload all database tables into individual output files from the AIJ backup file.
10. Using either the trigger technique or an application program, update the tables in the new database with the change data.
11. Start an on-line quiet point backup of the new database

12. Change logical names or the environment to specify the new database root file as the production database.
13. Restart the application on the new database.

Depending on the amount of application database activity, steps four, five, and six can be repeated to limit the amount of data that needs to be applied (and the amount of downtime required) during steps eight, nine, and ten, which are the final AIJ backup and apply stages.

Availability

The entire Oracle Rdb team is excited about the benefits that LogMiner will offer our customers. You can evaluate this feature today: LogMiner is currently in field test. We anticipate the initial production release early in the year 2000. Future Rdb ECOs and MUPs will continue to provide enhancements for the Oracle Rdb LogMiner feature.

Please feel free to join our field test program. Oracle Rdb version 7.0.3.1 is a required base software release for the Oracle Rdb LogMiner field test. Please use the Oracle Rdb Web site to sign up for the Oracle Rdb LogMiner field test and download the field test software.