



Oracle Rdb LogMiner and JCC LogMiner Loader

Bill Gettys / Philippe Vigier
Oracle Rdb Engineering

Based on presentations by
Jeffrey S.Jalbert, JCC Consulting, Inc.

Copyright 2002-4 JCC Consulting, Inc. and Oracle Corporation
used by permission

ORACLE




Why Data Replication is Important

Information access is increasingly important; disks
are increasingly cheap

- Ad hoc, reporting access interferes with OLTP
- Access to information needs to be continuous, but
 - Databases must sometimes be restructured
 - Databases must sometimes be isolated
 - Databases and systems sometimes fail
- Information must be protected from disaster
- Oracle Rdb is not always the right database management system


ORACLE

2




6 Methods

Replication Method	Type	Year
Journal Backup/Restore Method	Journal	~1980
Replication Option for Rdb	Table	~1985
Hot Standby	Journal	~1995
LogMiner/Loader	Journal	2000-2
Application Based	Table	
Shadowing/Mirroring	Disk	~1990




3



CLM/LML

- **Transactional**
 - One or many source transactions = one target transaction
- **Event driven or scheduled (Static LogMiner)**
- **Excellent performance on source database**
 - Uses journals, not tables
 - Takes advantage of hardware disk cache; no database hot spot
- **Excellent performance on target database**
 - Multiple load threads now supported
- **Multiple target databases supported**



4



CLM/LML (Cont.)

- No Geographic Limit
- Low network overhead
- Write your own loader if you like
- Lots of flexibility
 - Logical data model
 - Physical implementation
 - Supports Rdb, Oracle, Tuxedo targets; more possible
 - Read/write access to target possible (be careful)

ORACLE

5



What is the Rdb LogMiner?

- Public access to AIJ file contents
 - Oracle does not publish internal AIJ format
 - Command line interface
 - Use AIJ contents without performing recovery
- Extract changed table information from AIJ backup files
- Offers very high performance, low overhead
- Extracts changes to file, mailbox, pipe or your callable routine

ORACLE

6



What Rdb LogMiner Does

- The Rdb LogMiner reads the AIJs to produce a stream of database changes.
- The LogMiner can run in either of two modes.
 - Static
 - Runs on backup copies of AIJ files (must start with quiet point)
 - Can process AIJs on the same or a different machine.
 - Continuous
 - “Near Real Time” (needs to wait for the commit on the source).
 - Runs on a machine where the database is open.

ORACLE

7



How Rdb LogMiner Works

- Most information is already in the journal
 - Additional information stored - ignored by recovery
 - System relations used to retrieve metadata
- Capturing changes from AIJ means minimal impact to production database
 - Single, quick metadata retrieval
 - AIJ size will grow
 - Content of deleted rows is now written to the AIJ
 - CLM process usually reads cached data just written to the AIJ

ORACLE

8



What Isn't Rdb LogMiner?

- Not a Data Loader
 - The data extracted from AIJ backup files can have several targets: file, mailbox, pipe, callback
 - None of them are a database
- Data not in the source table format
 - Additional information is added to the record by LogMiner

ORACLE

9



Extract Operation

- All specified tables extracted at once (one pass through AIJs)
- Multiple output streams are possible
 - One output stream per table or multiple tables delivered to one output stream
 - Skip some tables
- One or more input AIJ files (processed sequentially)
- Output record order within transaction not predictable
- Only committed transactions are output
 - Output in transaction commit sequence order

ORACLE

10



What Does LogMiner Output?

- Data from user tables
 - Inserts and modifies
 - ‘Final’ record content per transaction
 - Deletes
 - Record content just prior to delete
- Only Committed transactions
 - Rolled-back transactions ignored
- Final form of row returned
 - Row modified many times in the same transaction results in a single modify output record
 - Transaction that inserts and deletes a row in one transaction results in only a delete record being extracted

ORACLE

11

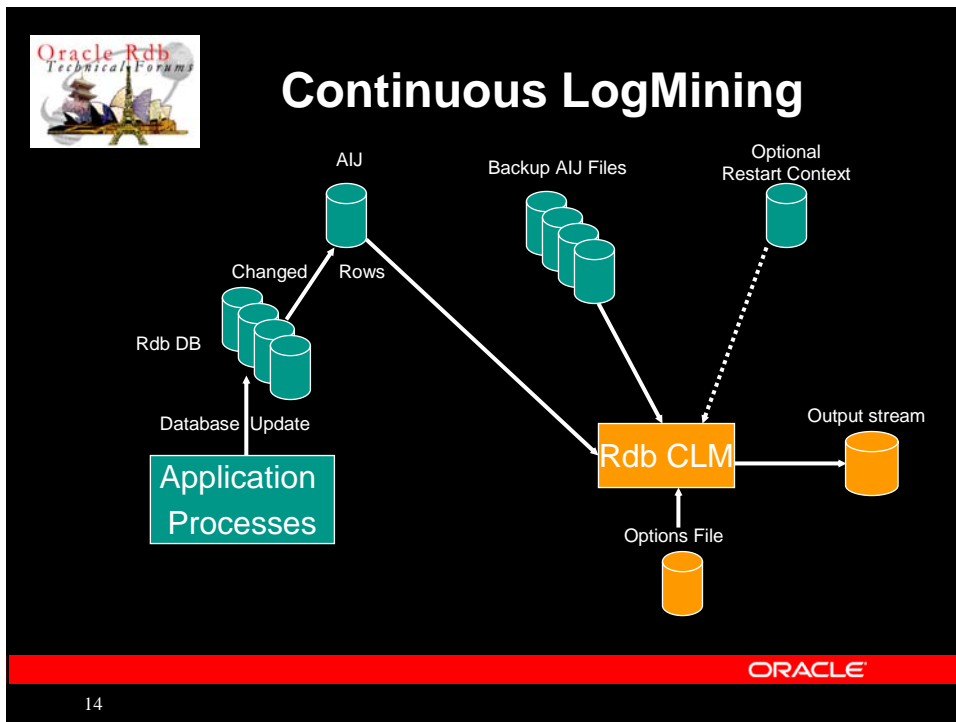
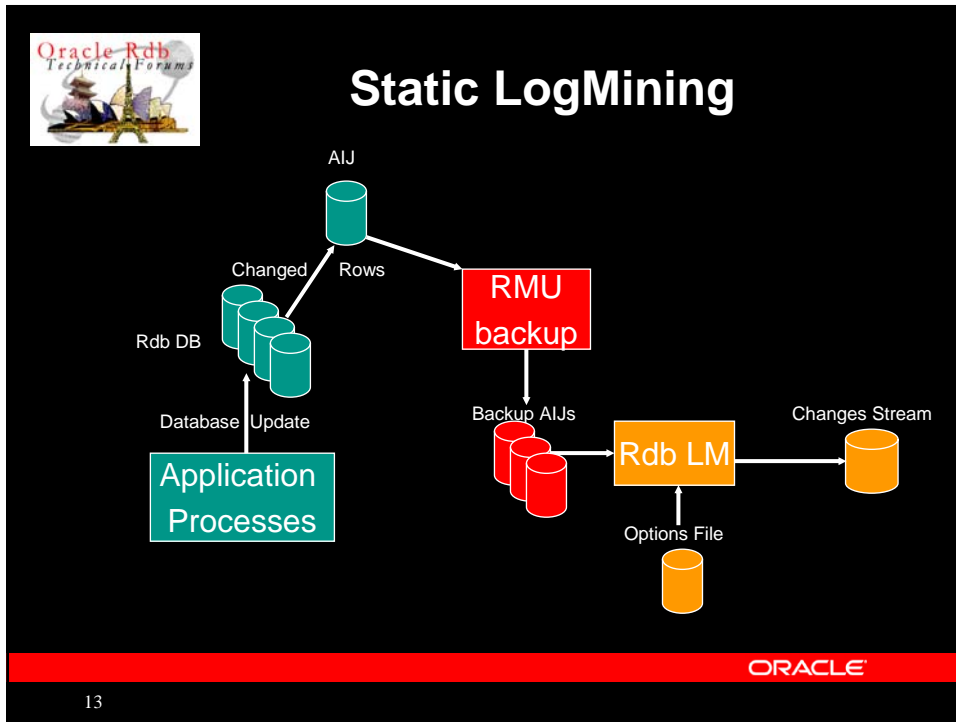


What Does LogMiner Output?

ACTION	“M” or “D”
TABLE_NAME	table name string
RECORD_ID	relation ID
RECORD_LEN	length of data record
NBV_LEN	null bit vector count
LDBK	logical DbKey
START_TAD	date/time of transaction start
COMMIT_TAD	date/time of transaction commit
TSN	transaction sequence number
RECORD_VER	metadata version of record
RECORD data	record contents
RECORD_NBV	null bit indicator array (1 bit per field)

ORACLE

12





Enabling LogMiner

- Requires offline access

```
$ RMU/SET LOGMINER/ENABLE MY_DB
or
$ RMU/SET LOGMINER/ENABLE/CONTINUOUS MY_DB
```

- Requires AIJ to be enabled
- Continuous LogMiner requires circular AIJ
- Backup database and AIJ after enabling

ORACLE

15



LogMiner Options File

- One file per table or the same file for multiple tables
- Options file example:

```
table=account                ,output=account_output_file.dat
table=account_audit          ,output=audit_output_file.dat
table=account_batch_processing ,output=account_output_file.dat
table=account_contract       ,output=account_output_file.dat
```

- Output types:
 - File
 - Mailbox
 - Pipe
 - Callback routine

ORACLE

16



Examples (Static LogMiner)

```
$ RMU/UNLOAD/AFTER_SOURCE.RDB SOURCE.AIJCK -
  /TABLE=(NAME=CUST,OUTPUT=CUST.DAT, -
    RECORD-DEFINITION = RDB-LM-RPT-CUST.RRD)

$ RMU/LOAD TARGET.RDB RDB-LM-RPT-CUST CUST.DAT -
  /RECORD-DEFINITION = FILE = RDB-LM-RPT-CUST.RRD -
  /COMMIT_EVERY = 1000
```

- Or Replace the table qualifier with an Options file

```
$ TYPE TABLES.OPTIONS
TABLE=MYTBL, OUTPUT=MYTBL.DAT
TABLE=SALES, OUTPUT=SALES.DAT
$ RMU/UNLOAD/AFTER_SOURCE.RDB SOURCE.AIJCK -
  /OPTIONS = FILE = TABLES.OPTIONS
```

ORACLE

17



Examples (Continuous LogMiner)

- This example demonstrates unloading three tables (EMPLOYEES, SALES, and CUSTOMERS) to a single mailbox.
- In order to include transaction commit information, the /Include=Action=Commit qualifier is specified:

```
$ DEFINE MBX$ LOADER_MBX:X
$ RMU /UNLOAD /AFTER_JOURNAL /CONTINUOUS MFP.RDB -
  /INCLUDE = ACTION = COMMIT -
  /TABLE = (NAME = EMPLOYEES, OUTPUT = MBX$: ) -
  /TABLE = (NAME = SALES, OUTPUT = MBX$: ) -
  /TABLE = (NAME = CUSTOMERS, OUTPUT = MBX$: )
```

ORACLE

18



Limitations

- Segmented strings, computed by field, temporary tables not captured
- Optimized AIJ files cannot be used
- Truncate or Drop table not captured
- Tables using vertical record partitioning (VRP)
- Changes to a table when the table definition has changed within an after-image journal file.
- Do not use the OpenVMS Alpha High Performance Sort/Merge utility

ORACLE

19




Loading

What do you do with the resulting data?

- Purchase a tool that does the work for you
 - JCC LogMiner Loader
- Write your own code
 - This requires understanding of all the unusual things that happen in the process. (see JCC LML documentation at <http://www.jcc.com/>)

ORACLE

20

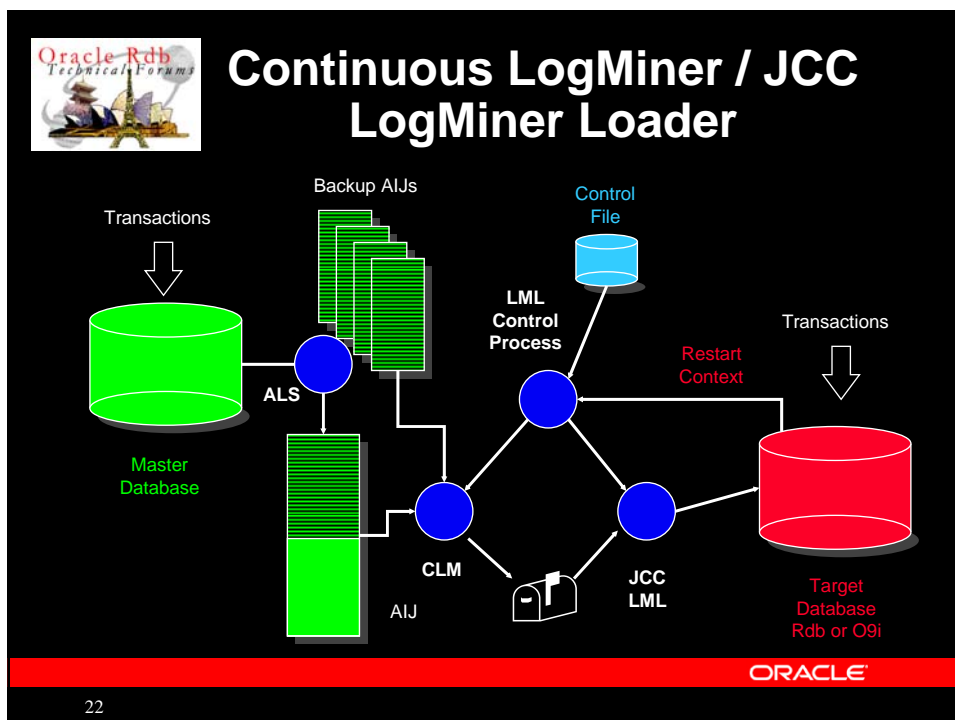


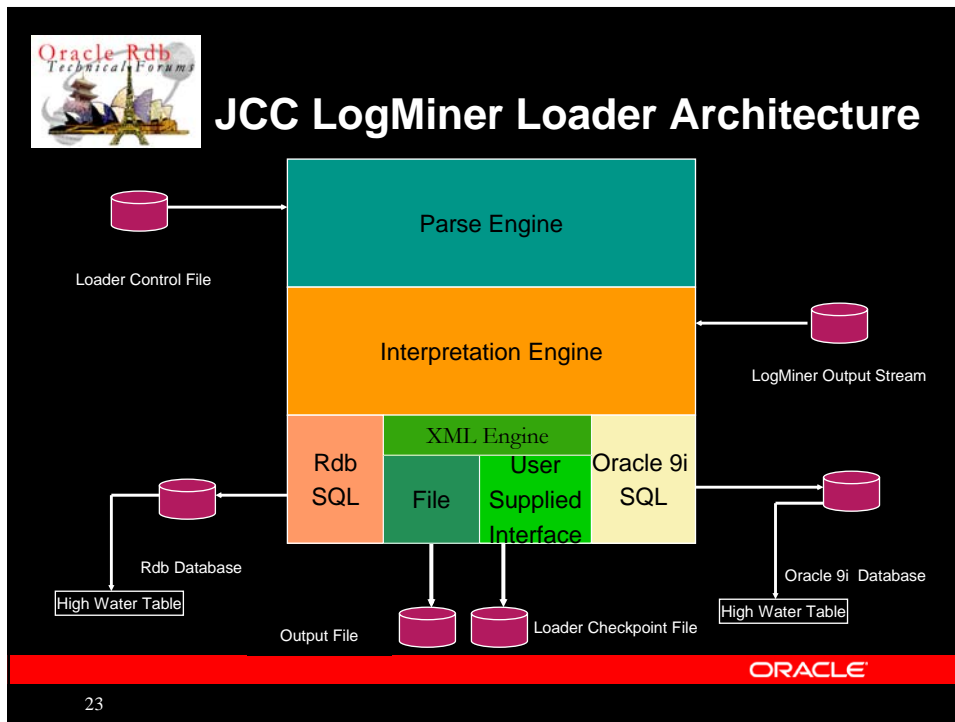
JCC LogMiner Loader Overview

- The Loader
 - Reads the binary files or output stream produced by the Rdb LogMiner and records the changes in a target data store
 - Reads a control file that specifies actions for the session
- The target data store may be one of the following:
 - Rdb database.
 - Oracle 9i database.
 - User-supplied API (XML document per source transaction.)
 - Disk file.

ORACLE

21





23

JCC LogMiner Loader Overview

- The Loader does its work within a recoverable and restartable context.
 - Maps one or more source database transactions into a single output transaction.
 - Stores checkpoint data at transaction boundaries in either the target database or in a local checkpoint file.
 - LogMiner AERCP
 - TSN of last transaction committed to target
- The Loader's actions are specified in a special Loader control file.
 - Type of target.
 - Tables and columns to transmit to the target.
 - Columns to materialize in the target, e.g. "Action."

ORACLE

24



LogMiner Loader Control File

- Control file controls all characteristics of replication
- Most database definitions will be long and complex.
 - LML Automatically generates the metadata portion of the Loader control file with supplied utility.
 - “Guesses” primary keys.
 - Must regenerate when the database metadata is changed.

ORACLE

25



Sample of Control File

```

!
LOADERNAME~DEMO
LOGGING~STATISTICS~COMMIT,RUNTIME,TIMER
!
! Table name, max version, what to do with table
TABLE~ACCOUNT_STATUS_CODE_TABLE~3~Replicate
!
PRIMARY KEY~ACCOUNT_STATUS_CODE_TABLE~ACCOUNT_STATUS_CODE~1~4~0~14~0
COLUMN~ACCOUNT_STATUS_CODE_TABLE~OLD_ACCOUNT_STATUS_CODE~2~1~0~14~0
COLUMN~ACCOUNT_STATUS_CODE_TABLE~CODE_TEXT~3~20~0~14~0
COLUMN~ACCOUNT_STATUS_CODE_TABLE~CURRENTLY_USED_FLAG~4~1~0~14~0
!
TABLE~ADDRESS_TYPE~3~Replicate
PRIMARY KEY~ADDRESS_TYPE~ADDRESS_TYPE_CODE~1~4~0~14~0
COLUMN~ADDRESS_TYPE~CODE_TEXT~2~20~0~14~0
COLUMN~ADDRESS_TYPE~USER_SELECTABLE_FLAG~3~1~0~14~0

```

ORACLE

26



Identifying Rows in Target Database

- The Loader can
 - Replicate rows to a target database
 - Write change records for audit or other purposes
- Replication produces tables exactly like the source table
 - Insert
 - Update
 - Delete
- Requires a primary key that does not change
 - Target database may be very different from the source database
- For tables without keys, the Loader can materialize a pseudo-key column, `originating_dbkey`
 - Guaranteed to be unique.
 - DBKeys can be reused so source table reorganization is more complicated.

ORACLE

27



Checkpoints

- The loader applies its changes in transactions
- The Loader retains restart context in the target database.
 - Requires a new table
 - The TSN of the last transaction committed
 - The AERCP
- For non-database output the Loader writes to a special checkpoint file
- Multiple source database transactions may be grouped in a single target database transaction
 - Checkpoints still work correctly

ORACLE

28



Near Real Time

- The LogMiner does not emit any records for a transaction until it finds a commit record
- All records for that transaction are emitted in a single group
- Therefore the LogMiner and Loader will always be “behind” the activities in the production database

ORACLE

29



The CLML Control Program

- Reads the target database or local checkpoint file to obtain restart information
- Creates a sub-process to run the Loader
- Creates a sub-process to run the continuous LogMiner
- Monitors the health of the two sub-processes
- Acts as a logging sink for the log files for the two sub processes (uses more mailboxes for this)
- Sends out operator alerts if something goes wrong
- Listens for shutdown request
- Drains mailbox during shutdown

ORACLE

30



Example Running Static Logloader

```
$ define jcc_logminer_loader_input <unload file>
$ define jcc_logminer_loader_init <control file name>
$ define jcc_logminer_loader_db <target database file name>
$ jcc_logminer_loader [-c] [-override override_type] -
  [-sequence seq_opt] -
  [-i input_name [-it input_type]] -
  [-o output_name] [-ot output_type] -
  ["-I" initialization_file_name]
```

Can also be defined in control file by:
output-rdb-synch-target_db-record

ORACLE

31



Running Continuous Loader

```
$ DEFINE JCC_ADD_CLM_SHARED_READ TRUE
$ DEFINE JCC_LOGMINER_LOADER_HW_RESPONSE CREATE|QUIT
$ DEFINE JCC_AIJ_backup_spec <directory name>
$ DEFINE JCC_ADD_CLM_LOG TRUE
$ DEFINE JCC_ADD_CLM_TRACE TRUE
$ DEFINE JCC_ADD_CLM_DEBUG TRUE
$ @jcc_tool_com:jcc_run_clm_lml.com <source database> -
  <LogMiner options file> -
  <LogMiner Loader control file> -
  [<restart override tag> <LogMiner restart context> -
  <Loader sequence number>]
```

ORACLE

32



Shutting down Continuous Loader

- In general, you should not terminate Loader or LogMiner jobs manually.
- For predetermined shutdown, use the procedure supplied with the JCC LogMiner Loader kit:

```
$ @JCC_TOOL_COM:JCC_CLML_SHUTDOWN.COM -  
  <loadername>
```

ORACLE

33



Obtaining Evaluation Licenses

- 90-day evaluation licenses are available
- Fill out form on JCC web page
- Distribution kit and documentation are available on the JCC FTP site:
 - <ftp://ftp.jcc.com/outgoing/JCCLogLoader/>
 - Requires license key to run
- Questions/information on Loader:
 - jeff@jcc.com

ORACLE

34



Customer Experience # 1

- U. S. Car Rental Business
- The internet really did change everything:
 - Web based travel sites replaced human travel agents
 - Far more rate requests
 - Far more rate changes
- Little capital available because of terror attack fears

ORACLE

35



Customer Experience # 1

- Database with normal transaction rate of 500 TPS.
 - About 2% of them update transactions.
- Loader database used for read only transactions in ACMS application.
- Target database is Rdb database.
- Maintains synchronization with one loader thread within a few (0-5) seconds.
 - Required commit interval of 50 because of bursts of do-nothing read-write transactions on night time batch runs.
- Average transaction duration reduced by 50%
- Solution is highly scalable


ORACLE

36



Customer Experience # 1

- Solution:
 - Replicate rate data to copy database
 - Retain master database for rate updates and reservations
 - Use Row Cache in both databases for speed (< 0.02 sec trans time)
 - Target databases never more than 5 seconds behind



ORACLE

37



Customer Experience #2

- Huge, multi database OLTP system
- Customer desires single Operational Data Store
- Source databases generate up to 5,000 change records per second during batch processing
- Replication target is Oracle 9i RAC database
- CLM-LML replication to target never falls more than 5 minutes behind

ORACLE

38



Customer Experience #3

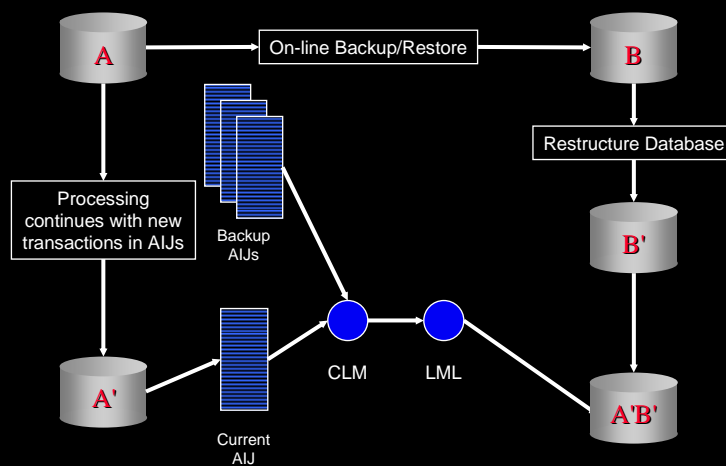
- Reorganize critical on-line 50Gb database in near-on-line mode
- Requirement: less than 1 hour downtime
 - Export/Import timed at 20 hours
 - Unload/Reload timed at 12 hours
- Solution:
 - Reorganize off line.
 - Apply AIJ iteratively to catch up.
 - One last AIJ application with database down.
 - Minimal down time.
- Database now has architecture to support growth.
- Jobs run in 1/2 of the time.

ORACLE

39



Customer Experience #3



ORACLE

40



Other Uses:

- Populate a reporting database
- Consolidate information from multiple systems to deliver information on the web
- Security audit changes to a table
- Populate a data warehouse
- Recover accidentally deleted records
-

ORACLE

41

COMMENTS ?
Q U E S T I O N S
&
A N S W E R S

ORACLE

ORACLE®

www.oracle.com/rdb

Philippe.vigier@oracle.com