



With No Production Rdb on NT What Do You Do Now?

(Creating Distributed Reference Databases)

**Jeffrey S. Jalbert,
Cheryl P. Jalbert,
Thomas H Musson,
Keith W. Hare**

JCC Consulting, Inc

Abstract

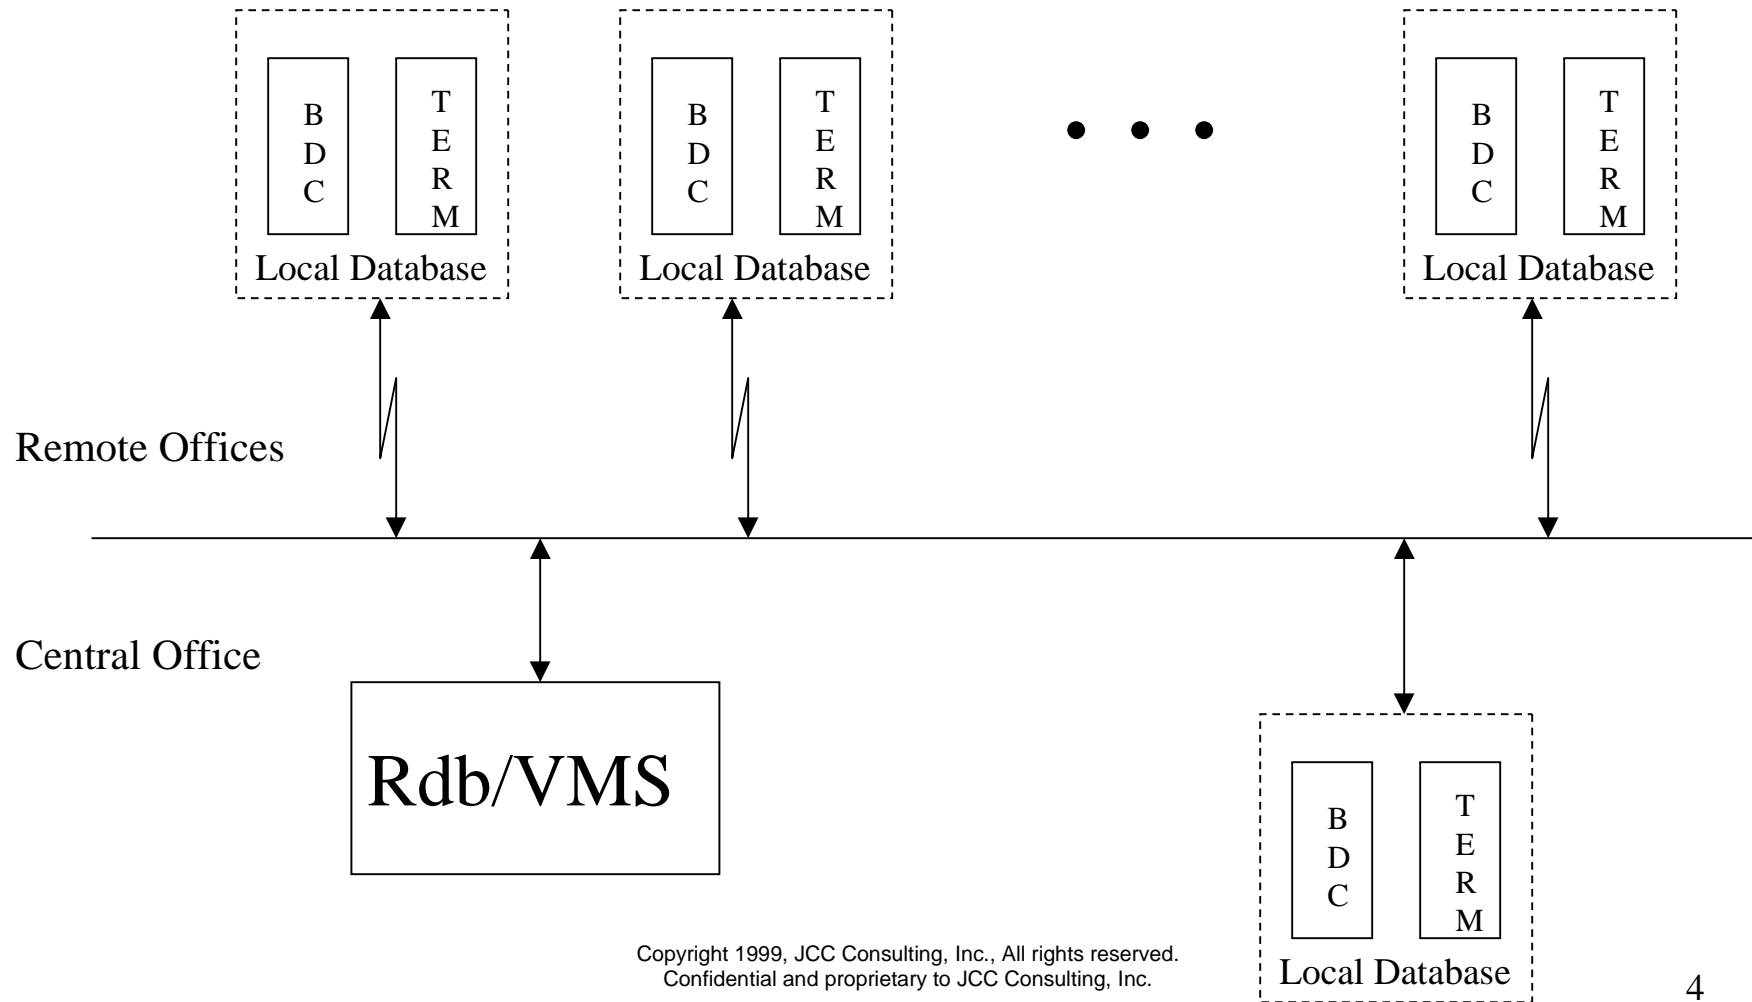
Rdb on Windows/NT promised us the ability to transport a number of applications using the Oracle Developer tools to an NT platform. Our plans had been to use these databases to provide backup reference databases should communications with a primary production database fail.

This presentation will outline how we intend to use the new characteristics of the log miner together with locally developed tools to maintain a number of reference Oracle 8 databases.

Computing Environment

- Enterprise has a central Rdb database on a VMS server
- Also supports a single NT domain for “PC” stuff
- Small remote offices are connected via 56 Kb lines
 - NT Backup Domain Controller
 - NT Terminal Server Edition
 - Hosts client GUI application

Computing Environment



Business Environment

- Enterprise is truly customer focused.
- Update transaction rates are modest
 - Aij grows between 25 Mb & 100 Mb per day, depending on day
 - Most of this is in indexes

Business Environment

- Computing support mixes queries and data update.
 - Lots of on-line queries for customer service
 - Modest on-line update rates
 - Frequent [every third day] large batch updates
- In a pinch,
 - Users can make hand notes and enter data later.
 - Users have a major crisis over not having query information available 100% of the time.
 - This is a real 24 x 7 operation [public utility]

Business Environment

- Unavailability of information can be a safety risk, up to and including death.
- Communications lines of all types have gone down at critical times.
- In fact, some of the most critical times [storms] specifically threaten the lines.
- In non-critical times, the central database does not have to be 24 X 7, *if* a reference database is available.

Problem Statement

- Provide local, supported, read-only access to production data should communications lines fail
- Retain lower-cost NT computing environment
- Provide central reference database while production db is down for maintenance
- Provide experience base for migrating away from VMS

Problem Statement

- Data must be up to date
 - Synchronized daily is sufficient
- One set of code that talks to local & remote databases
- Automatically detect network failure & change to local database

Some Considerations

- We briefly investigated using the “Replication Option” for distributed databases
 - Basically this is the old data distributor
 - But going through DBI
- Decided that it was too many layers
- DDAL has concurrency issues relating to the tables describing rows to be replicated
- Puts responsibility for updates on VMS server
- Rejected this idea

Steps to Replication

- Create an initial empty “clone database”
- Populate with data synchronized to a point in time
- Synchronize with master database periodically

Architectural Solution – Defining the Metadata

- Create Oracle 8i databases on NT systems in each office
 - Includes central office for availability reasons
 - Serves as the template for all remote databases
 - Can move data at LAN speeds
 - Table definitions identical to production database
 - Identical indexing, but that is not required

Architectural Solution – Defining the Metadata

- Scripts can be generated by reading Rdb schema with SQL procedures
 - Maybe with RMU/EXTRACT=noDomain??
- Will not support referential integrity in slave databases
 - Will allow us to perform synchronization in any order

Architectural Solution – Populating the Base Data

- Export base production data via automatically generated scripts
 - Use SQL to process system tables to create views, procedures and data files containing base data suitable for loading via Oracle loader
 - Exploring unloading Rdb database directly with native Oracle tools
- Load resulting data into empty Oracle 8 tables

Architectural Solution – Synchronizing Data Changes

- Use the Rdb Log Miner to extract changes to the production database
 - Explored using audit tables but that is too costly in terms of additional tables
- Send resulting changes across the link
- Use stored procedures & knowledge of primary keys to maintain remote data
 - Only one message over slow network links per remote row updated/deleted/inserted
- Much of this is what we would do to support a data warehouse

Log Miner Condensation

- Results so far indicate that running the log miner will result in a reduction factor of between 5 and 10 on the AIJ size
 - Most of the AIJ is comprised of indexes
- Will preprocess on a per-table basis to ensure that we send only the latest version of each row
- This results in between 5 to 20 Mb of data to transfer nightly
 - Can accomplish in two hours or so at 11:00 PM
 - Leaves plenty of time to maintain remote database

Technical Problems to Solve

- Manage the time duration to slave the first database. We have not yet measured that.
 - We don't have a good feeling for O8 index build rates yet
- Retrofitting primary key constraints on the production database.
 - Not really a problem to solve, but rather a task to do
 - Most tables without primary key definitions are audit tables, but required for the application
 - Not a problem because audit tables are insert-only for this application

Technical Problems to Solve

- Metadata changes will require even more careful maintenance
- Require quiet point AIJ backups before & after change
- Exploring ways to avoid this slow down of software deliveries

Technical Problems to Solve

- Managing the execution of the update code across several systems.
 - Synchronize schedule on NT to schedule on production VMS server
- Bandwidth needs for propagating the daily updates
 - These must be validated. They are estimated only at this point

Technical Problems to Solve

- Quality of Oracle 8 optimizer vs. the Rdb dynamic optimizer

Non Problems

- SQL scripts to generate the unload data and load into Oracle 8 already exist from other work
- Creating the stored procedures will be straightforward because we can parse the existing Rdb metadata.

Non Problems

- User passwords will be synchronized across VMS & NT
 - VMS will use NT authentication
- Application mediates all login processing so it knows user's passwords to both NT & VMS
- Ad-hoc reports are increasingly becoming based on Developer Reports and performed in background by NT

Non Problems

- The Oracle 8 physical database design will be abstracted from the existing Rdb design
 - Pretty much is one table/index per storage area
 - We will need to see how Oracle 8 stands up to hundreds of table spaces.
 - The database servers will have plenty of memory to support the Oracle database

Non Problems

- Managing automatic switchover of client code
 - Detect pipe failure to production database by trapping `dbms_error_code` (-3121,-3113,-3114)
 - Already do that so we can support users better when we experience SQL services failures
 - Reconnect to the fail-over database behind the scenes & alert the end-user
 - Database configurations are embedded in an application .INI file read at startup.
 - Set our security model such that all updates are now forbidden
 - Our security model is read from the database at application startup.

Business Problems

- Training users to understand that they can be in a fail-over environment
 - Legacy code is VMS and will not fail over.
 - With current version of forms we will “flash” the database on the screen title bar
 - In a later version we will be able to change the window background color
 - Ad-hoc reports will also be failed over.

Projected Implementation

- First remote office will be supported sometime in late Q1 or early Q2 2000
 - Depending on availability of Log Miner software
 - Also depending on our own personal bandwidth
- Remaining offices will be supported in Q2

Questions & Comments



Copyright 1999, JCC Consulting, Inc., All rights reserved.
Confidential and proprietary to JCC Consulting, Inc.