Dbvisit
THE SMART ALTERNATIVE

What a DBA needs to know about the differences between Oracle physical and logical Database replication

Presenter: Arjen Visser, Founder of Dbvisit Software

Introduction

Arjen Visser @dbvisit
Founder of Dbvisit Software Limited
Technical Entrepreneur with passion for Oracle Database

Past Experience:

- DBA
- Unix admin/project manager
- Datawarehouse developer/programmer
- Speaker at OOW, NZOUG, CLOUG, RMOUG, AOUG, Collaborate, DOAG

Dbvisit

Dbvisit
THE SMART ALTERNATIVE

- Dedicated to database replication software products
- Based in New Zealand with sales offices in US, Europe and Asia
- Trusted by 700+ Companies
- Worldwide leader in DR solutions for Oracle Standard Edition
- Regular presenters at Oracle events such as OOW and Collaborate
- Passionate about Oracle Technology…..its all we do!!

Dbvisit product lineup

Dbvisit
THE SMART ALTERNATIVE

World-leading Disaster Recovery Solution for the Oracle SE database

Affordable Alternative to GoldenGate: Data Replication for Oracle-sourced databases

Objective

1. Understanding the differences between logical and physical replication
2. Understanding the advantages and disadvantages of each
3. Understanding how they work and the underlying concepts
4. Understanding the use cases of each

Replication options

1. Trigger based - No longer used
   - Changes to source database are required
   - Performance impact
   - A lot more maintenance
2. Redo log based - Preferred method
   - No changes to source database
   - Lower impact to source environment
   - Less maintenance
   - DDL changes can be replicated

What is Redo?

Contains history of all database changes

1. Used for database recovery
2. Written from redo log buffer first and then to redo log (after commit, 3 seconds or 1/3 full)

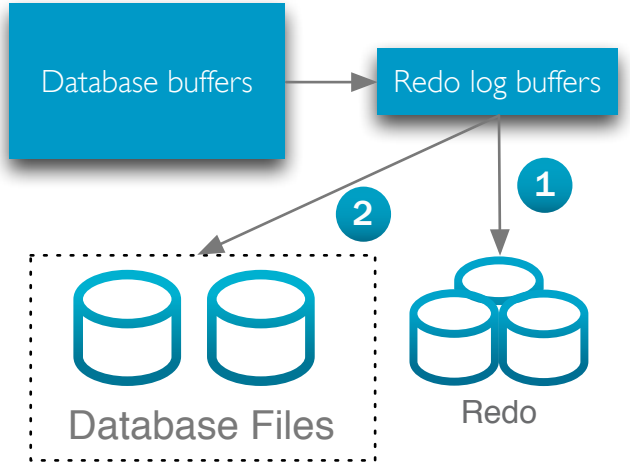

What is Redo?

Redo will include:

1. Changes made by Insert, Update, Delete, Merge (DML)
    - Change vectors. These describe the transition of a database block from one state to another.
2. DDL

Redo does not include:

1. Actual text from DML
2. Select statements

Redo records behaviour

- Redo gets the database change first
- Multiple change vectors grouped together create a redo record.
- Redo records transitions the database from one state to another state.
- All or none of the change vectors from a redo record will be applied. Transaction consistency (ACID)

Redo - Most important

- Does not contain SQL !
- Contains operations to database blocks
- There are 150 different redo operations
- They are in binary format specific to OS and endianness
- They are specific to the database SID and DBID
- They are specific to the Oracle version

Two Oracle Replication Types

Dbvisit
THE SMART ALTERNATIVE

Based on Oracle Redo:

Physical Replication

- Binary copy
- Data Guard

Logical Replication

- Logical copy
- GoldenGate

Oracle replication based on redo

Physical Replication

- Use complete redo blocks for replication
- Information is unaltered
- Target applies binary changes (redo blocks)
- Target is in constant "recovery" mode

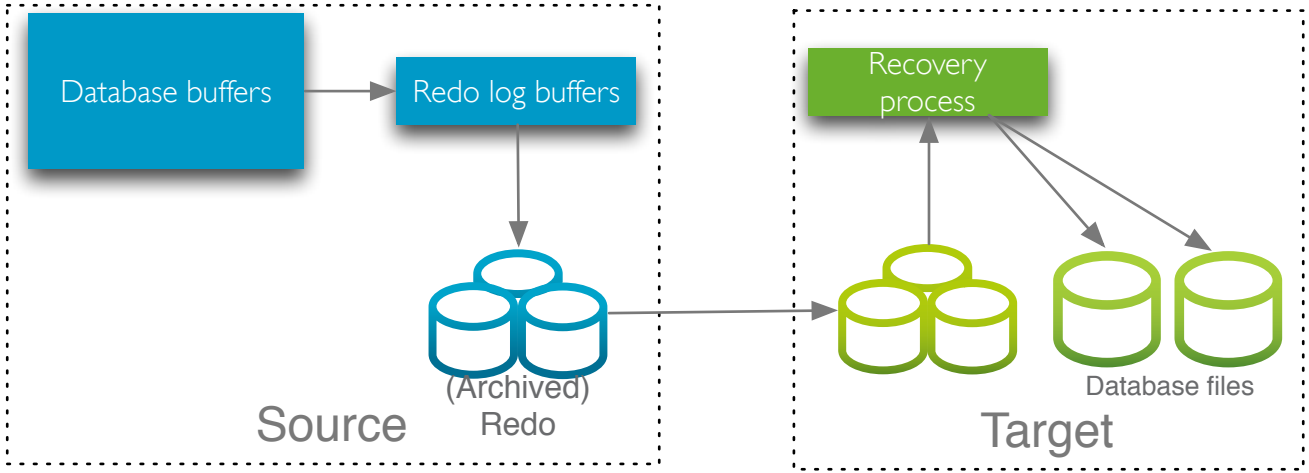

Oracle replication based on redo

Logical Replication

- Mining of Redo - Extract information from redo blocks for replication
- Information is altered/transformed into SQL
- SQL is applied at target
- Target is normal read/write database

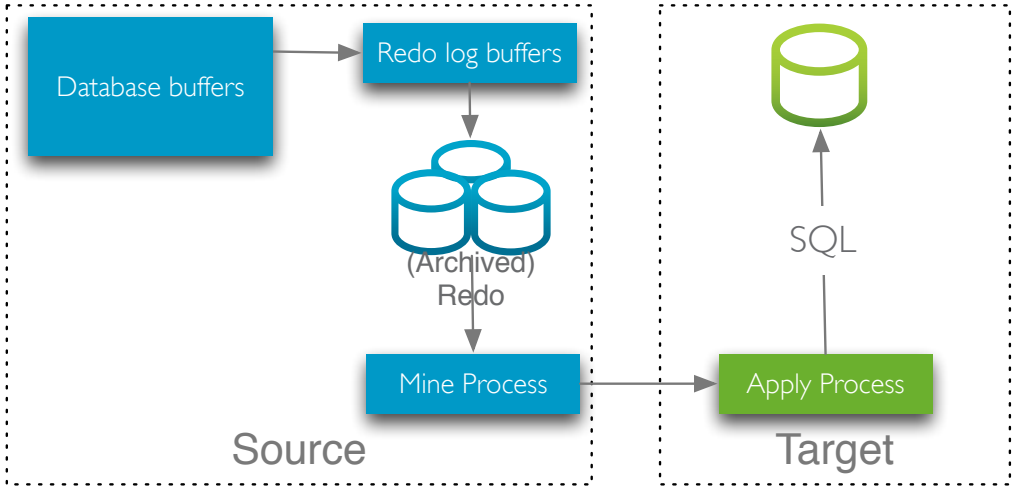

Two Oracle Replication Types compared

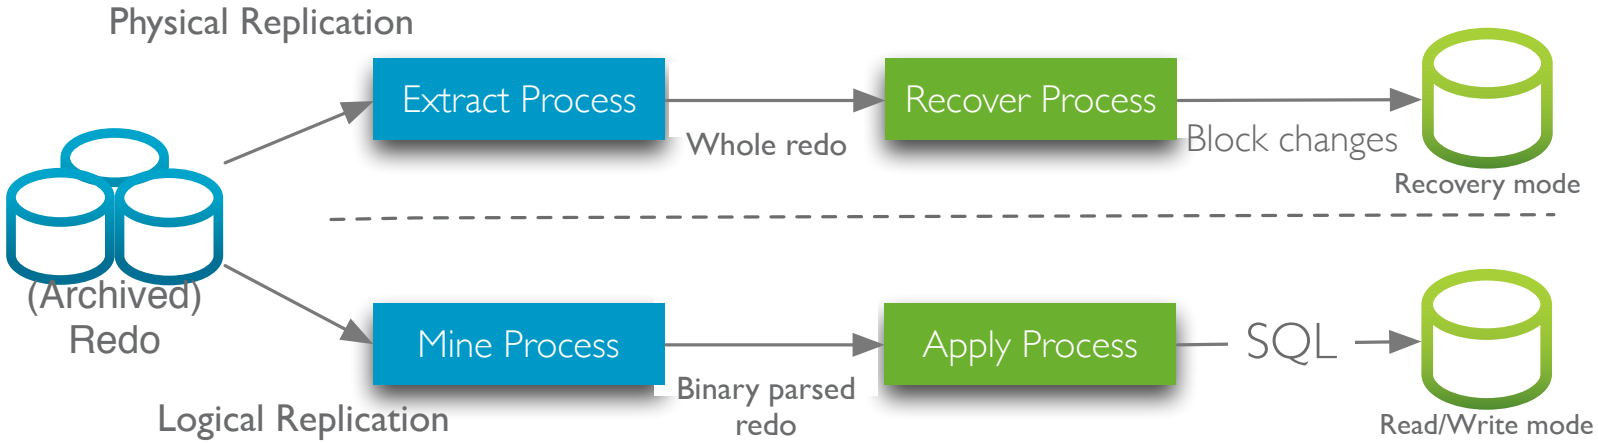

Two Oracle Replication Types

Dbvisit
THE SMART ALTERNATIVE

1 - Physical Replication

- Use database block level (redo) to keep up to date
- Physically the replicated database is same as source database
- 100% binary copy, structure and data
- Referred to as **primary** (source) and **standby database** (target)
- Standby database is in constant recovery mode
- Examples
  - Oracle Data Guard (Oracle Enterprise Edition only)
  - Dbvisit Standby (Oracle XE, SE1, SE, EE)

Physical Replication

Dbvisit
THE SMART ALTERNATIVE

Advantages

- 100% binary copy, structure and data
  - rowid, indexes, FK, security, DB links
- Offload backup to standby database
- Recover data files from standby database
- Applying changes to standby is fast and efficient
- Standby server can have fewer resources
- Easy to create standby database using rman or third party tool
- Easy to manage
- Low admin overhead
- Standby database can be used for read only queries
- Roll reversal or graceful switchover is possible

Physical Replication

Disadvantages

- Replication is all or nothing
- Same OS must be used
- Same Oracle version must be used
- When Standby database is in recovery mode it is unusable
- NOLOGGING changes break the standby database
- All Archives are needed, no single sequence can be skipped
- Activation of standby database breaks the link with primary
- When standby database is in read only mode, it cannot apply latest changes*

* Unless Active Data Guard is used

Two Oracle Replication Types

Dbvisit
THE SMART ALTERNATIVE

2 - Logical Replication

- Uses SQL statements to keep database up to date
- Replicated is database has the same data
- Separate physical database structure
- Examples:
  - Oracle GoldenGate
  - Streams
  - Dbvisit Replicate

Logical Replication

Dbvisit
THE SMART ALTERNATIVE

Advantages

- Subset of data can be replicated
- Across different OS and Oracle versions
- Target database is normal read-write database
- Different indexes and security policies can be employed
- Target database can be non-Oracle (Hadoop, MySQL MS SQL etc)

Logical Replication

Dbvisit
THE SMART ALTERNATIVE

Disadvantages

- Not as efficient as physical replication
- Need the same resources as source database
- SQL is applied which can cause potential conflicts and locks
- SQL applied is not the same as on source
- Database not the same structure (tablespaces, indexes, partitions, MV)
- Not everything is replicated (Data types, passwords, sequences, rowid)
- Requires data-instantiation (data is in sync before replication starts)

Database requirements

Physical replication

- Force logging is recommended

Logical replication

- Force logging is recommended
- Supplemental logging must be turned on

Any changes not recorded in redo will not be replicated !

- Direct load
- APPEND hint in SQL statement

```
INSERT /*+ APPEND */ INTO emp
  SELECT * FROM t_emp;
COMMIT;
```

Licensing

Physical replication

- Need a full Oracle license on standby
- Does not have to be the same **edition** as source*
- Does not have to be the same **metric** as source

Logical replication

- Need a full Oracle license on target
- Does not have to be the same **edition** as source
- Does not have to be the same **metric** as source

* EE features cannot be used if Standby is SE or SE1

Deeper dive

Logical Replication

How does redo log mining logical replication work

Dbvisit
THE SMART ALTERNATIVE

1. Redo logs are mined for LCR* as Oracle is writing to the logs (real time)
2. LCR are pieced together in correct order (ie row chaining, row migration, RAC)
3. Filtered LCR are written to a file or a queue
4. Transferred to target server
5. Translated to SQL statements
6. Run against target database

* Logical Change Record - describes a change to a row

Typical logical replication flow

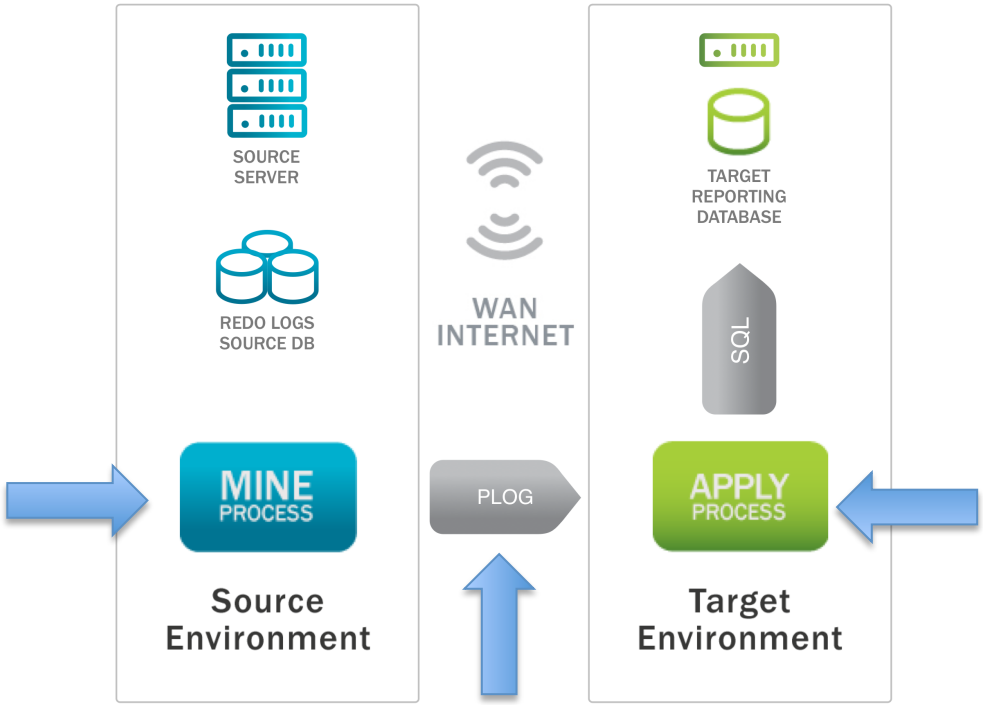

Mine

Mines the redo logs and converts into a parsed log

PLOG

Parsed logs – binary files specific to solution (trail file in OGG)

Platform independent

Apply

Converts parsed log data into target DB native SQL

Conversion process

Set based SQL operations on the source database are converted to row-by-row SQL changes on the target database.

Why?

- This is the way that Oracle writes to the Redo
- Redo does not contains SQL, it contains db block changes

Observations:

- SQL is not the same on source as on target
- True for all logical replication based solutions
- Each SQL on target only affects 1 row

Logical replication principle

Source	Target
`update PRICES set PRICE = PRICE - (PRICE * .10) where PRODUCT_CAT = 'OLD_STOCK';`	`update PRICES set PRICE = 10 where PROD_ID = 101;` `update PRICES set PRICE = 23 where PROD_ID = 102;`

Observations:

1. Source SQL updates 2 rows, then 2 individual update statements are produced
2. PK has been added to the WHERE to ensure row-by-row
3. Price formula has been replaced by hardcoded value

Considerations with logical based replication

Data instantiation

Before replication can start, data needs to be in sync between source and target. Ensure:

- All data is captured including in-flight data
- All required indexes
- PL/SQL

Use

1. Datapump (export/import)
2. Rman duplicate
3. Standby database (resetlogs)

Considerations with logical based replication

Commit strategies

1. Optimistic commit (do not wait for commit before mining and applying)
2. Wait for commit before applying (do not wait for commit before mining)
3. Wait for commit before mining and applying

Examples of products with above strategies:

1. Dbvisit Replicate
2. Streams
3. Golden Gate

Considerations with logical based replication

Dbvisit
THE SMART ALTERNATIVE

Generic

1. Replicated data is same order as source (based on SCN)
2. Transaction integrity is maintained (replication is serial)
3. Be-aware of triggers on target database
4. Atomic transactions are maintained (commit & rollback)
5. Conflicts pauses the whole replication to ensure integrity
6. Conflicts warn of data divergence

Deeper dive

Physical Replication

Standby protection modes

There are three protection modes for the primary database with physical replication:

Maximum Performance: Transactions on the primary commit as soon as redo information has been written to the online redo log. Transfer of redo information to the standby server is asynchronous, so it does not impact on performance of the primary.
Available in SE & SE1

Maximum Availability: Transactions on the primary do not commit until redo information has been written to the online redo log and the standby redo logs of at least one standby location.
Not available in SE & SE1

Maximum Protection: Transactions on the primary do not commit until redo information has been written to the online redo log and the standby redo logs of at least one standby location. If not suitable standby location is available, the primary database shuts down.
Not available in SE & SE1

Keeping the standby database up to date

Dbvisit
THE SMART ALTERNATIVE

Standby databases are held up to date by applying archive (redo) log files

- Standby database is in constant recovery mode
- For Oracle EE, online redo (using standby redo) or archive logs can be used
- For Oracle SE and SE1, only archive logs are used (no standby redo logs)
- Uses standby controlfile
- Oracle EE

```
sqlplus> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE DISCONNECT FROM SESSION;
```

- Oracle SE and SE1

```
sqlplus> recover standby database
```

- Archive log files from ALL threads are needed with RAC

Recover RAC Standby Database on SE and SE1

```
SQL> recover standby database
ORA-00279: change 54284704 generated at 05/08/2012 10:07:48 needed for thread 1
ORA-00289: suggestion : +DATA
ORA-15173: entry 'ARCHIVELOG' does not exist in directory 'ORCL'
ORA-00280: change 54284704 for thread 1 is in sequence #2326

Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
/oracle/oraarch/ORCL/thread_1_seq_2326.1592.782677117     ORA-00279:
generated at 05/08/2012 10:07:48 needed for thread 2
ORA-00289: suggestion : +DATAORA-00280: change 54284704 for thread 2 is in sequence #2714

Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
/oracle/oraarch/ORCL/thread_2_seq_2714.1599.782677117
ORA-00279: change 54292400 generated at 05/08/2012 10:18:35 needed for thread 2
ORA-00289: suggestion : +DATA
ORA-00280: change 54292400 for thread 2 is in sequence #2715
ORA-00278: log file '/oracle/oraarch/ORCL/thread_2_seq_2714.1599.782677117' nolonger needed
for this recovery

Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
Cancel
Media recovery cancelled.
```

Give name of actual location of archive log file

Now it wants thread 2 seq 2714

Thread 2 seq 2715 is not available yet, so cancel

Cancel media recovery

Considerations with physical based replication

1. Failover should not be treated lightly
2. Archive missing or NOLOGGING invalidates standby database
3. Database patching requires standby database consideration
4. Requires regular DR testing

Replication Use cases

Dbvisit
THE SMART ALTERNATIVE

Disaster Recovery

Physical Replication

Logical Replication

Data Distribution HA

Logical Replication

Zero downtime Migration

Logical Replication

Reporting Databases BI

Logical Replication

Physical Replication

Trusted in 80+ countries . . .

. . . By 700+ companies.

Dbvisit
THE SMART ALTERNATIVE

Thank you
Questions and Answers?

info@dbvisit.com
@dbvisit