Logical Replication in 12cR2: What are the Options Now?

16. November 2016

Vit Spinka | Chief Architect, Dbvisit

Agenda

Dbvisit
move it, use it

- General overview
- Options obsolete in 12*c*
- What is still available
- Oracle GoldenGate
- Oracle Data Guard: Logical Standby
- Third party options

Vit Spinka

Dbvisit
move it, use it

- Working with Oracle Database since 8i
- 12c Oracle Certified Master
- Oracle ACE Associate
- Principal developer of Dbvisit Replicate
- … which gets its data by parsing Oracle redo logs

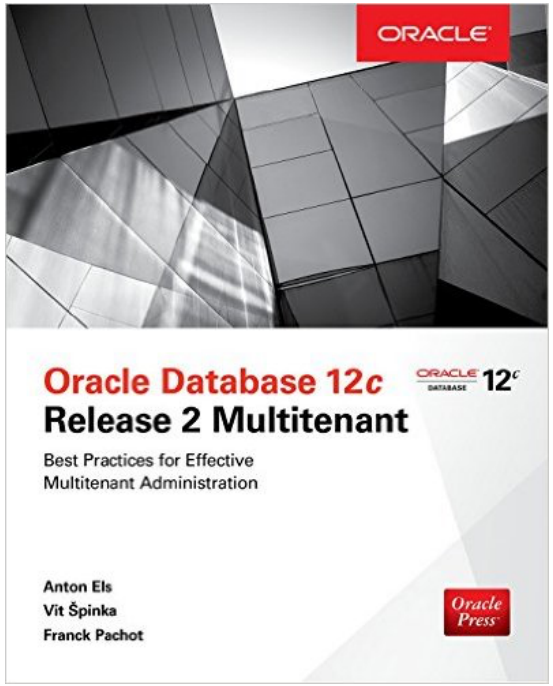

- @vitspinka
- vit.spinka@dbvisit.com
- http://vitspinka.com, http://vitspinka.blogspot.com

General Overview

- We are not talking about physical replication
- Physical standby still works (and replicates the whole CDB)
- And storage-based replications still work, too
- We look at logical replication
- We are interested in data as seen by the users
- This is much more complex than physical standby…

Logical replication principles

Dbvisit
move it, use it

- We work with logical changes
- Translating events into one-by-one changes
  - Statement-level replication is not much of a thing in Oracle
- We use PK/UK to identify rows
  - Unlike physical, which essentially uses rowid
  - Can work without PK/UK, but you really don't want duplicate rows
- Possibility of conflicts
  - Data not as expected (when applying changes)
  - Or simply don't add up (when seeing it as stream of changes)

Basic approaches

- Triggers
  - After every change, trigger runs and gathers :old and :new
  - Slower, synchronous, affecting user sessions
  - Issues with DDL changes
  - Can be done internally, using faster code path
- Materialized view logs
  - Special case of triggers

Basic approaches

- Redo, LogMiner
  - Reading changes from redo
  - More difficult to parse, e.g. row chaining/migration, trailing NULLs
    - Only modified values are put into redo
    - Plus supplemental logging
  - But has every (recoverable) change
  - So can handle even non-SQL changes, e.g. partial LOB updates

Targets

Dbvisit
move it, use it

- Can be a 1:1 copy of the source table on target
  - Or a subset
- Target is a relational database
- Changes are replicated as one-line SQL
- Conflicts can occur
  - **This is good**: it tells us about data issues, bugs, updates at target
  - We thus see the issues and can work with them
- We work with the data on target as we would on source, it still looks the same

Targets

- Can be a simple stream of events, stored in a table
- Or pushed out
  - Message queues (JMS)
  - Special API (XStream)
  - Out-of-box supported non-relational targets (Kafka)
- We work with the data differently than on source
  - We see stream of changes
  - We act upon events
  - Or we comb through complete history

Mandatory note about DG and ADG

- Logical replication is not DR
  - It's not reliable enough
  - It's not simple enough
  - Just use physical standby – it is built for DR, it's easy to understand at 2 a.m. when the building is on fire
- Even so, if all you want is to offload reports to a different database:
  - Consider Active Data Guard
  - It's still physical standby, easy to use and robust
  - But can be used for read-only queries
  - Additional licensing required, though

Options obsolete in 12c

Dbvisit
move it, use it

Oracle obsoleted or deprecated a couple of replication options

Reasons are partly technical – there were a lot of them, and multitenant brought changes to internal structures

Perhaps more important is pricing – these options were free (included in Oracle Database license)

But the suggested replacement is not…

Oracle CDC

Dbvisit
move it, use it

- Stores changes as records in a staging table in local database
- Thus often it needs some further processing by other tools
  - CDC/Events become a hot topic nowadays… but too late for this feature
- Can be synchronous, using triggers
  - SE or EE database
- Can be asynchronous (reading redo)
  - EE (but still no extra cost on top of EE)

A simple example

Dbvisit
move it, use it

OPERATION$	EMPNO	ENAME	SAL
I	7839	KING	5000
UO	7839	KING	5000
UN	7839	QUEEN	6000
D	7839	QUEEN	6000

Oracle CDC

- Internally used by some replication tools
- Deprecated in 11*g*
- Obsolete and removed in 12*c*

Oracle Streams

Dbvisit
move it, use it

- Asynchronous, reading redo logs
- Introduced in Oracle 9*i*R2
- Probably the most used replication option
  - Although with its own deal of issues, e.g. large transaction handling
  - And difficult to recover after things break down
- Free
  - Included in EE license

Oracle Streams

Dbvisit
move it, use it

- However, deprecated in 12*c*
- Still works
  - But no new 12*c* features are supported
  - Which first and foremost include multitenant
  - They simply decided not to do the changes necessary for multitenant
- No development on Streams anymore
  - Instead the code was used in Oracle GoldenGate internal capture

Oracle Streams

Dbvisit
move it, use it

- Streams had many powerful features
  - Rulesets
  - Transformations
- Basic use case is maintaining 1:1 copy of source tables, but rules/filters can customize it
- Oracle provides tool for conversion of (simple) replications to OGG:
  - Streams2OGG (1912338.1)

Advanced Replication

Dbvisit
move it, use it

- Updateable materialized views
- Based on materialized views
  - And materialized view logs
  - Works nicely in disconnected environments
  - Simple 1:1 copy of source – or even multimaster
- Refreshed on demand (or on schedule)
  - So not really real-time (but can be close, depending on volume)
- Or can be synchronous
  - But that is fragile, it's essentially distributed commit

Advanced Replication

- Deprecated in 12*c*
  - *ORA-26949 Advanced Replication is not allowed in a container database.*
- This does not affect the “usual” materialized views, though
  - And these can be remote (Basic Replication)

Oracle GoldenGate

Dbvisit
move it, use it

This is the Oracle-envisioned one-and-all tool for logical replication

Powerful

And expensive

Oracle GoldenGate

Dbvisit
move it, use it

- Oracle acquired GoldenGate in 2009
- Supports many sources
  - Most development is done for Oracle source
- And many targets
- Classic capture
  - Uses to read redo logs and parse them
  - The “old” way
- Integrated capture
  - Uses database API, internally based on Streams code

Oracle GoldenGate

Dbvisit
move it, use it

- Integrated capture is the future
- That's really where the development is done
  - This is the main synergy of Oracle acquiring GoldenGate
- Better datatype support, faster
- But needs 11.2.0.3 (or .4), as it relies on code changes inside the database

- Similarly, there is integrated and non-integrated replicat
  - Non-integrated issues SQL using OCI
  - Integrated uses database API

Oracle GoldenGate

Dbvisit
move it, use it

- Current market leader
- A strong player
  - If you need the best performance
  - If you need the best datatype support
  - And don't care about user-friendliness
  - And don't care about price
  - And don't care about speed of support
- Paid add-on option: management pack for Enterprise Manager
- Check carefully licensing, esp. if not's the basic Oracle-Oracle

Oracle GoldenGate Multitenant Support

Dbvisit
move it, use it

- Multitenant is supported
- But only with integrated capture
- Can use one capture for all PDBs
  - And one replicate for each target PDB
  - Still consider using multiple captures, if the replications are logically independent
- As with any new features
  - Test your use cases, e.g. hot cloning, relocation, common objects

XStream API

Dbvisit
move it, use it

- Integrated capture and replicat use an API..
- And Oracle made it public
- XStream In
- XStream Out
- Requires Oracle GoldenGate license
  - And since that implies you have full Oracle GoldenGate, why bother with API?
- And C/Java programming

Oracle GoldenGate for Big Data

Dbvisit
move it, use it

- For non-RDBMS targets like
  - Flume
  - HDFS
  - Hive
  - Hbase
  - Kafka

Oracle Data Guard: Logical Standby

Logical standby was introduced a long time ago

… still, not many people use it

But Oracle sees future for it, and it’s alive and well

Data Guard Logical Standby

Dbvisit
move it, use it

- Part of Data Guard
  - Part of EE license
- Uses the same transfer technology and setup
- Intended use is for whole database (can skip tables/schemas)
- But internally, it's similar to Streams
  - Similar issues: needs unique rows, limited datatype support

Logical Standby, an easier way to Streams

Dbvisit
move it, use it

- Not as versatile as Streams
- Supposed to be more “set up and forget” than Streams
- But at the same time, this limits its usefulness
- As it’s created from physical standby, similar same platform/version rules apply

Status in 12c

- Still officially supported
- Not many limitations
  - But for example, 12.2 application containers are not supported
- When talking to people at Oracle
  - It is not seen as the future of general-purpose replication – and instead Golden Gate is recommended
  - .... but this comes with a price tag

Rolling upgrade

Dbvisit
move it, use it

- The real reason for keeping logical standby…
- And where the development is done…
- Near-zero downtime migrations
  - But based on logical standby, so still needs same OS, same endian
  - Although keep an eye of expanding list of cross-platform support: 1085687.1 (which is more restrictive than 413484.1 for physical)
- Convert physical to logical, do the upgrade, convert back

DBMS_ROLLING

- For a single primary/standby, it's easy to do an upgrade/migration manually
- But if you have more of them, you want to always have the primary protected by a physical standby… then use DBMS_ROLLING
- First upgrade half of the standbys
- Then switch over
- Then upgrade the other half
- Then clean up and revert back to original config (convert back to physical)

Third party options

Oracle is not the only player in logical replication

There are other options

Generally cheaper, more user-friendly

Even if with more limitations

Might take time to pick up 12.2 support

Let’s have a look at two of them

Dell SharePlex

Dbvisit
move it, use it

- Dell acquired it with Quest in 2012
  - Now Dell sells it again to Francisco Partners
  - A private equity company, thus it will be sold again, soon
- Supports multitenant
  - Each PDB has its own configuration, capture and apply process
- Only Oracle as a source
- Uses its own redo parser
- Supports Kafka, Change history

Dbvisit Replicate

Dbvisit
move it, use it

- Uses its own redo parser
- Only Oracle as a source
- Supports multitenant
  - Each PDB has its own mine and apply process
- Supports Kafka, Event Streaming

Dbvisit
move it, use it