

Active Data Guard

Svetozár Kapusta

Distributed Database Operations
Workshop


November 27th, 2009

- Introduction
- Data Guard Concepts
- Active Data Guard
- Real Time Query Performance
- 11G2 ADG Features
- Conclusions

- Data Guard plays a key role in the MAA best practices
- Data Guard automatically maintains standby DB(s) as transactionally consistent copy(ies) of the primary DB by
 - transmitting primary DB redo data to the standby DB(s)
 - applying the redo data to the standby DB(s)
- If the primary fails, the standby DB can be quickly activated
- Minimizes downtime for both planned and unplanned outages
- Primary and standby communicate over Oracle Net Services

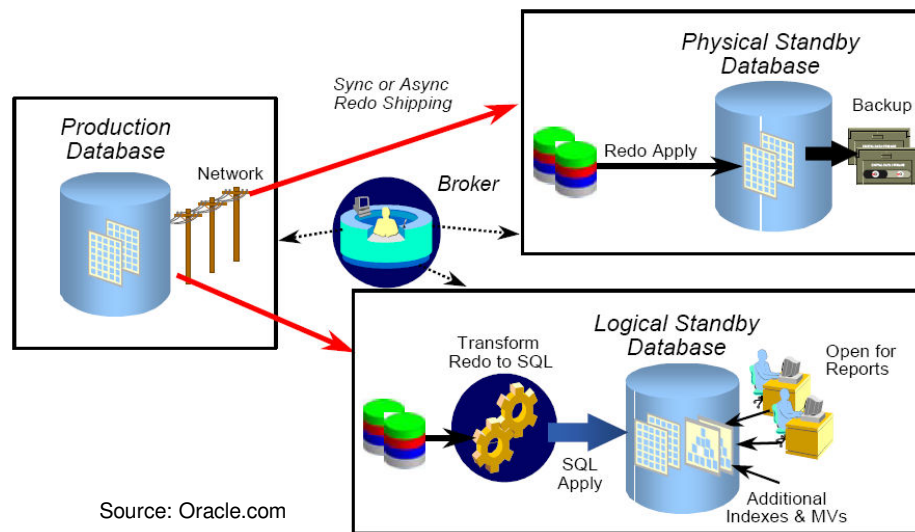
- Introduction
- **Data Guard Concepts**
- Active Data Guard
- Real Time Query Performance
- 11G2 ADG Features
- Conclusions

- 2 types of standby databases:
 - Physical standby
 - uses Redo Apply to maintain a block for block, exact replica of the primary database.
 - Logical standby
 - uses SQL Apply and contains the same logical information as the primary database, although the physical organization and structure of the data can be different.



Source: Oracle.com

- Physical standby DB
 - Can be used to offload the primary DB of the overhead of performing backups, since it is an exact replica
- Logical standby database
 - additional flexibility of being open read-write
 - additional local tables can be added to the DB
 - local index structures can be created to optimize reporting or to utilize the standby DB as a data warehouse



Source: Oracle.com

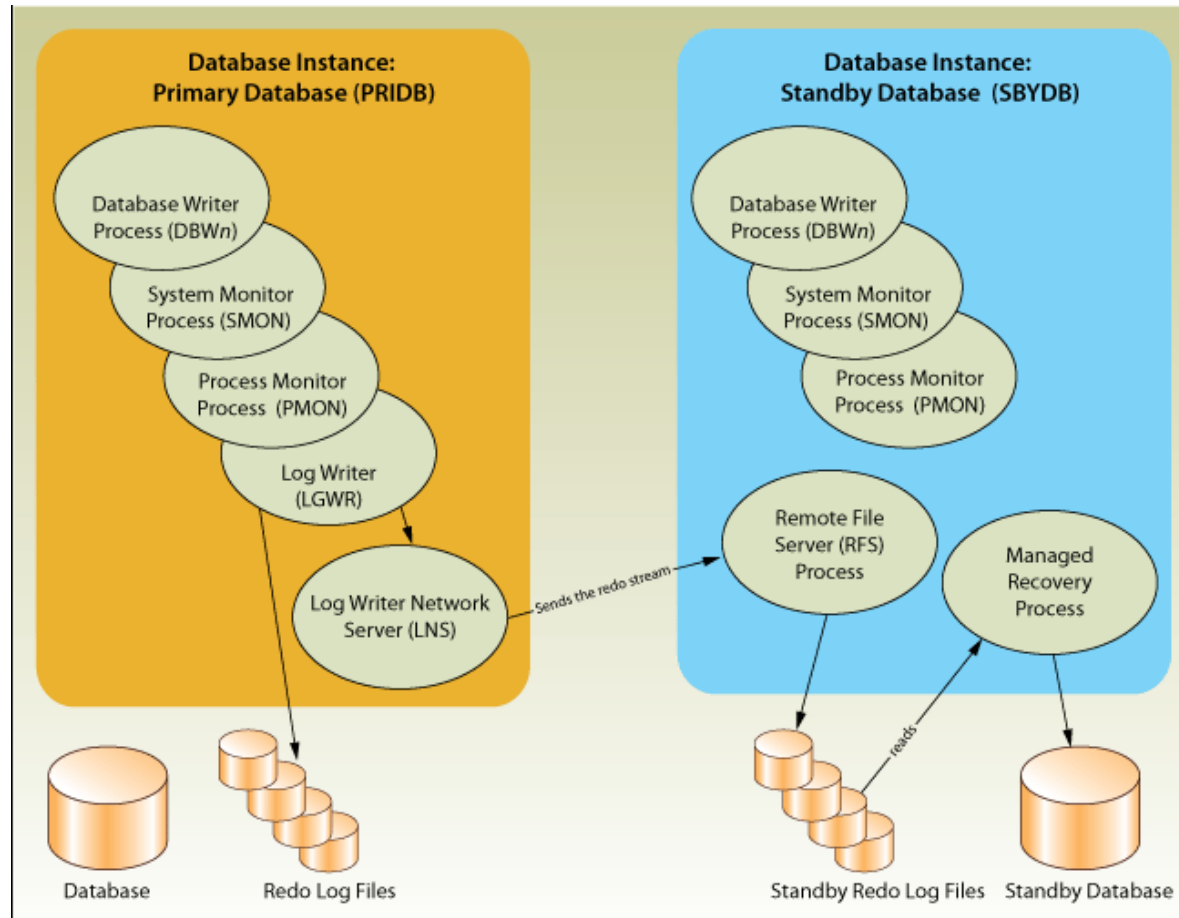
- Synchronous redo transport (SYNC)
 - Primary DB waits for confirmation from the standby that redo has been written to disk before it acknowledges commit success to the application
 - Provides zero data loss protection
 - Primary DB performance affected by time required for the standby redo log file I/O to complete and network round-trip time
- Asynchronous redo transport (ASYNC)
 - Primary DB acknowledges commit success to the application without waiting for acknowledgment that redo has been received by the standby DB
 - Almost no performance impact on primary DB
 - Potential for a small amount of data loss

- Maximum Protection (SYNC)
 - Zero data loss, Double failure protection
 - Stall primary DB until acknowledgement is received from the standby DB
- Maximum Availability (SYNC)
 - Zero data loss, Single failure protection
 - Stall primary DB until acknowledgement is received or NET_TIMEOUT threshold period expires – then resume processing
- Maximum Performance (ASync)
 - Potential for minimal data loss
 - Primary DB never waits for standby acknowledgment

- Switchover is a planned, zero data loss operation
- Guarantees standby data can not be modified independent of primary transactions
 - No split brain
 - No data corruptions
- Redo Apply
 - Primary drains redo pipe
 - Standby applies all redo
 - Flips a bit in the control file
 - Changes role to primary
 - Opens standby as primary
 - No bounce required
 - Done
- SQL Apply
 - Primary drains redo pipe
 - Standby applies all redo
 - Flips a bit in the control file
 - Turns off the Guard
 - Changes role to Primary
 - Done

- Introduction
- Data Guard Concepts
- **Active Data Guard**
- Real Time Query Performance
- 11G2 ADG Features
- Conclusions

- Physical standby DB can be opened read-only while simultaneously receiving updates from the primary



- Easy to enable:
- `SQL> alter database recover managed standby database cancel;`
- `SQL> alter database open read only;`
- `SQL> alter database recover managed standby database using current logfile disconnect;`
- Write access possible via network links:
- `SQL> create database link write_testADG connect to sveto identified by test using 'test11g2';`
- `SQL> insert into test@write_testADG values (1,12);`
- 1 row created.
- `SQL> commit;`
- Commit complete.

- Manage primary and standby(s) from the same interface
- Simplifies failover and switchover operations
- In 11G2
 - `SQL> alter database open read only;`
 - Broker will jump in and automatically stop Redo Apply and the restart it after the open has completed
 - At switchover, if Active Data Guard in use at the target standby the original primary will be opened when it becomes a standby after the switchover

• 11G1

- SCN or V\$DATAGUARD_STATS to calculate lag
- `SQL> SELECT name, value, datum_time, time_computed FROM V$DATAGUARD_STATS WHERE name like 'apply lag';`

- NAME VALUE DATUM_TIME TIME_COMPUTED

- -----

- apply lag +00 00:00:00 10/25/2009 13:14:11 10/25/2009 13:14:11

• 11R2 view V\$STANDBY_EVENT_HISTOGRAM

- `SQL> SELECT * FROM V$STANDBY_EVENT_HISTOGRAM WHERE NAME = 'apply lag' AND COUNT > 0;`

- NAME TIME UNIT COUNT LAST_TIME_UPDATED

- -----

- apply lag 0 seconds 48612 10/25/2009 13:20:02

- apply lag 1 seconds 102 10/25/2009 13:15:09

- apply lag 2 seconds 16 10/25/2009 12:20:58

- apply lag 3 seconds 4 10/25/2009 11:15:56

- Introduction
- Data Guard Concepts
- Active Data Guard
- **Real Time Query Performance**
- 11G2 ADG Features
- Conclusions


- HW:
 - Primary: 2 RAC nodes, 2 disk arrays
 - Standby: 2 RAC nodes, 2 disk arrays
- SW: RHEL 4, Oracle 11G1 and 11G2 beta 2
- Installed via RMAN's "duplicate target database for standby from active database"

```
rman target sys@{DB_NAME}_primary auxiliary sys@{DB_NAME}_standby
RUN {
  allocate channel prmy1 type disk;
  allocate channel prmy2 type disk;
  allocate auxiliary channel stby type disk;
  duplicate target database for standby from active database
    NOFILENAMECHECK spfile set control_files='/tmp/control_stdby.ctl'
    set db_unique_name='{DB_NAME}_stdby';
}
```

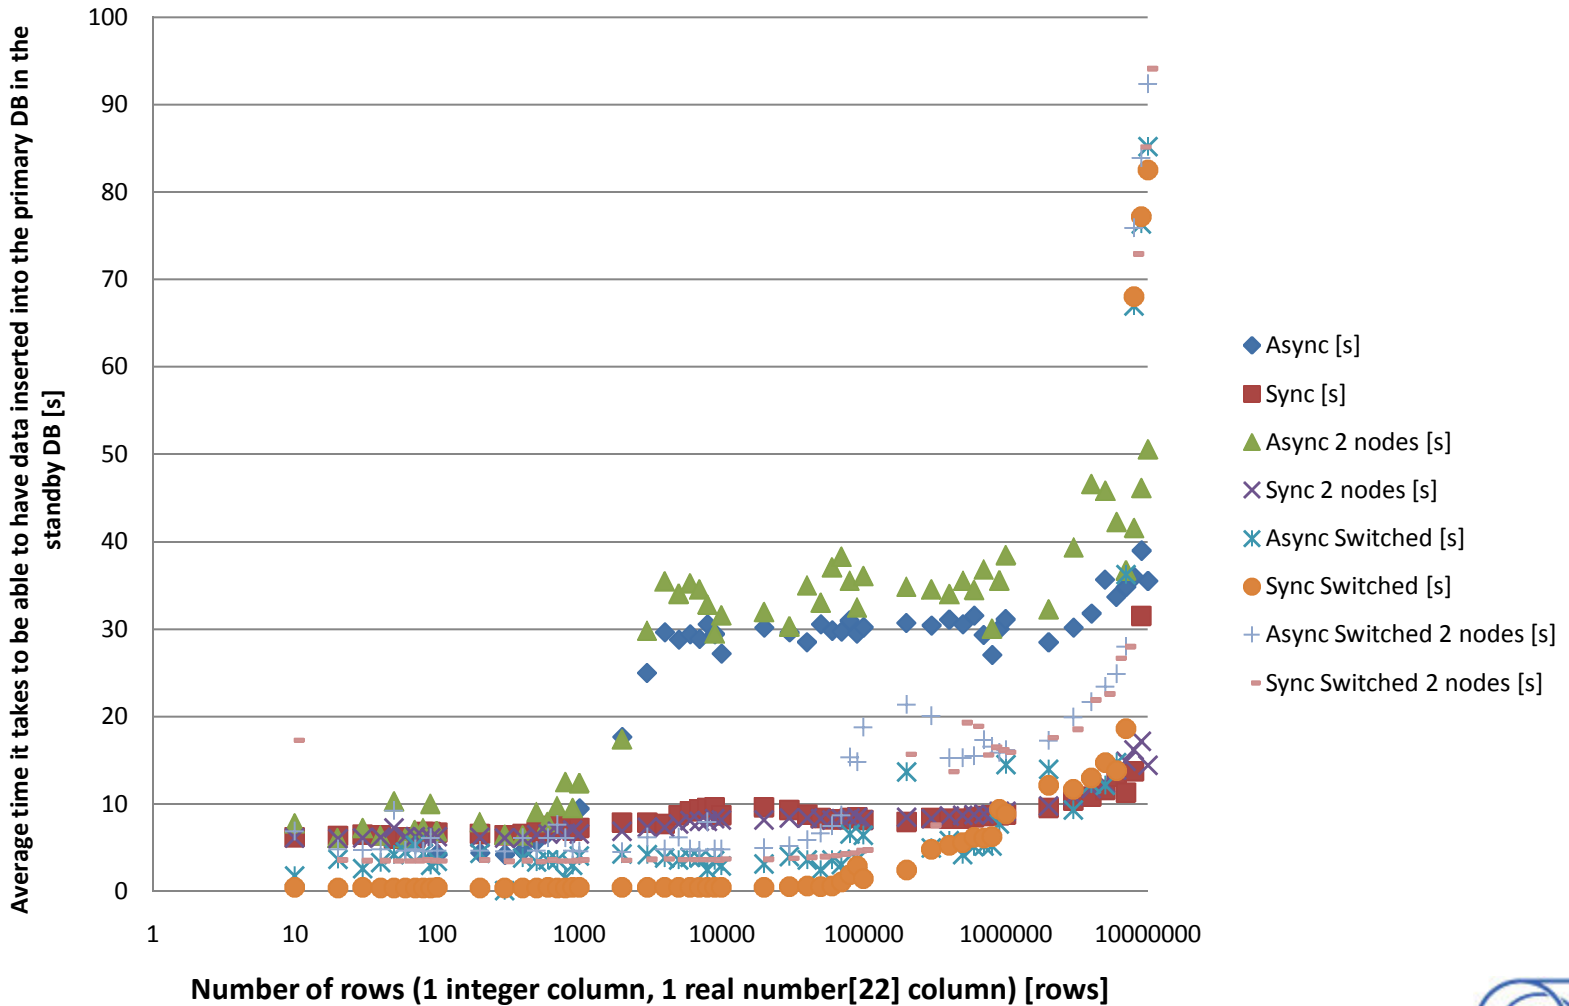
- RMAN 11G2 will try to continue where it left off

- Measuring algorithm:
- Loop X times over:
- Insert N data rows into the primary DB
- Measure the time it takes that any row appears in the standby DB
- Repeat the above for:
 - Varying the number of inserted rows in one transaction from 10 to 10e+7
 - 1 and 2 nodes of physical standby RAC active
 - Synchronous and asynchronous REDO transport
- Repeat all the above after a switchover

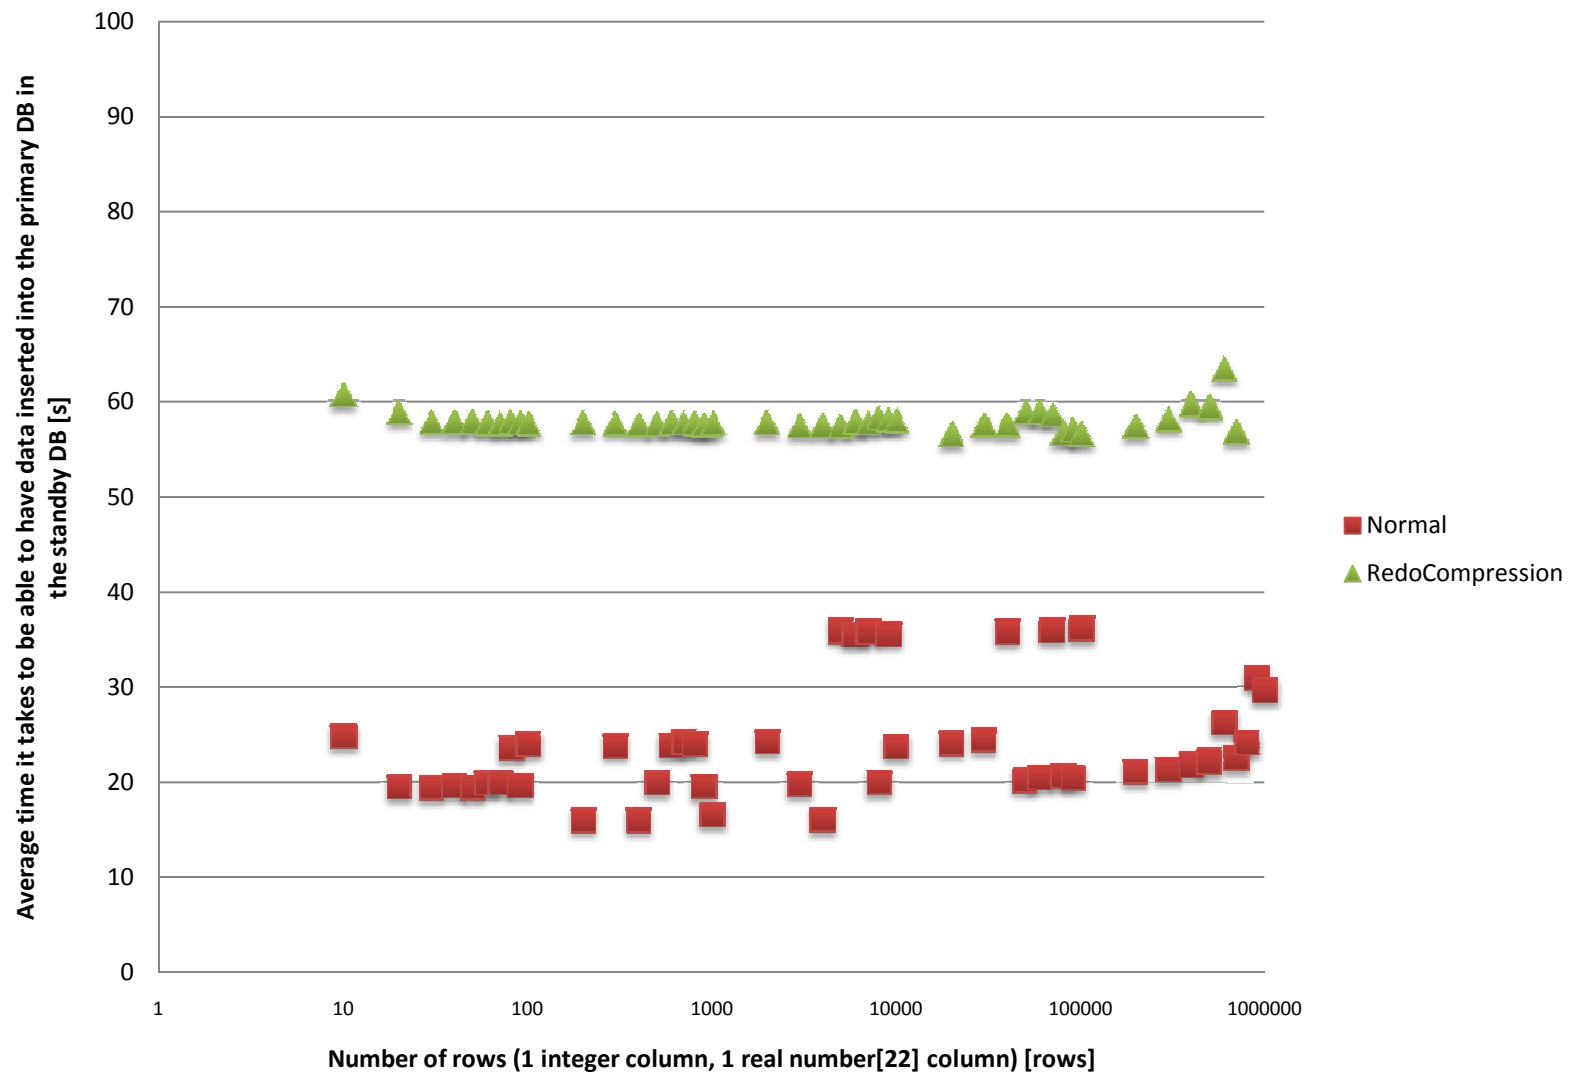

Active Dataguard Performance Using Real-Time Apply



Active Dataguard Performance Using Real-Time Apply



- 1 node standby RAC is slightly more performing then a 2 node one
- Synchronous REDO transport of course outperforms the asynchronous one in these tests
- Truncate table with a subsequent query on the table on standby gives ORA-08103 (normal behavior)
- Confirmed that the standby DB could be used for read-only at all times
- Verified the long term stability
- Performed a quick and smooth switchover



- Introduction
- Data Guard Concepts
- Active Data Guard
- Real Time Query Performance
- **11G2 ADG Features**
- Conclusions

- RMAN's fast incremental backups (11G1)
 - block change tracking enabled on physical standby
- Redo Transport
 - Support for up to 30 standby databases for a single primary database (was 9)
- Protection
 - un-sent redo in asynchronous configurations may be flushed to a standby before failover
 - zero data loss (if failed primary can be brought to mount state)
- Role Transitions
 - Redo Apply switchovers no longer require any standby instances to be shut down

- Physical standby DB to be open read-write for testing or any activity that requires a read-write replica of production data.
 - continues to receive updates generated by the primary, but doesn't apply them
 - applied to the standby DB automatically when the Snapshot Standby is converted back to a physical standby DB
 - Standby must have flashback feature enabled
 - **SQL> ALTER DATABASE CONVERT TO SNAPSHOT STANDBY;**
 - To roll back just restart in mount mode and:
 - **SQL> ALTER DATABASE CONVERT TO PHYSICAL STANDBY;**

- Standby DBs can be used to perform planned maintenance in a rolling fashion
 - Objective: upgrade primary and standby to new Oracle version, patchset, 32<->64 bit, Windows <-> Linux, single node -> RAC, migrating to ASM
 - Using SQL Apply
 - any upgrade from 10.1.0.3 onward
 - Using Redo Apply –Transient Logical Standby
 - any upgrade from 11.1.0.7 onward
 - Maintenance first performed on a standby DB
 - Production switched over to the standby DB when the maintenance tasks are completed
 - Only downtime is the time needed to effect a switchover operation

- New session setting **STANDBY_MAX_DATA_DELAY**
- NONE (Default) - queries will be executed regardless of the apply lag on that database
- Non-zero - queries will be executed only if the apply lag is less than or equal to **STANDBY_MAX_DATA_DELAY**
- If delay setting exceeded an error is returned
 - **ORA-03172: STANDBY_MAX_DATA_DELAY of 10 seconds exceeded**
 - Application then decides what to do
- Zero - queries guaranteed to return the exact same result as if the query were issued on the primary database otherwise the ORA-03172 is returned
 - Requires Maximum Availability and Real-Time Apply

- When Oracle discovers a corruption it marks the block as media corrupt, writes it to disk, and returns an error to the application:
 - **ORA-01578: ORACLE data block corrupted (file # 5, block # 140)**
- With ADG 11R2 the corrupted block on the Active Data Guard primary will be repaired from the standby (and vice versa)
- Alertlog:
 - **Requesting Auto BMR for (file# 5, block# 140)**
 - **Waiting Auto BMR response for (file# 5, block# 140)**
 - **Auto BMR successful**

- Active Data Guard is a very promising technology
 - High Availability
 - Disaster Recovery
 - Data Protection
 - High Performance
- Finally the standby can be used more than just waiting for the disaster
 - Open read only
 - Snapshot standby
 - Automatic repair of corrupt blocks
 - Block change tracing (RMAN fast incremental backup)
- DG easier to set up
 - RMAN's duplicate for standby from active
- We are looking forward to deploy it



DM

•Q&A

Thank you for your attention

