

Active Data Guard Oracle Open World - 2007

@Amazon.com

Grant McAlister
Principal Database Engineer

Uses of Active Data Guard

- Reporting copy of a primary system
 - Typically different queries than primary system
 - Gap may be in hours
 - (i.e. data from last night/last business day)
 - One copy
- Scalable read store
 - Subset of queries that run on primary
 - Gap in seconds/minutes - not hours
 - As many copies as needed to handle read load

Physical standby for reporting

- 8i
 - Destructive copy of database rolled forward to point in time.
 - Could use 3rd mirror to make this simpler
- 9i – 10g
 - Read-Only Mode – non destructive
- 11g
 - Real Time Query – Active Data Guard

Active Data Guard for Scalable Read

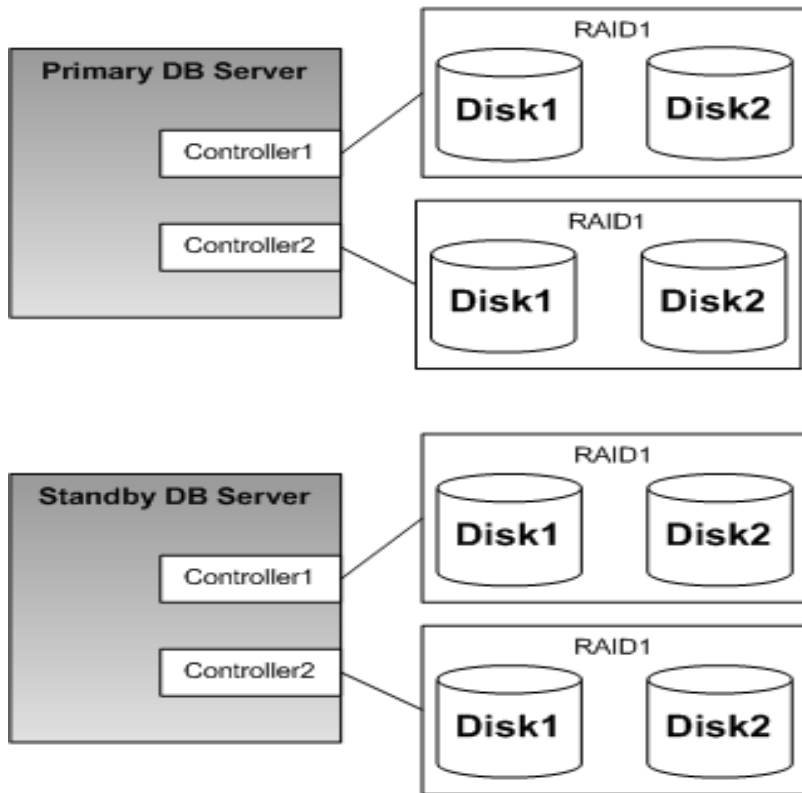
- Allow read only queries to scale beyond single db
- Higher availability for read only queries
- Can be configured to shed extra reads
- More efficient use of hardware

Data Guard Fast-Start Failover


- Ability to use less expensive equipment
 - Reduced need for multiple controllers for redo mirroring
 - Possible to depend on standby servers for protection
- Need 1 standby but using 2 is much better
 - Efficiency is <50% with 1 standby and <33% with 2
- How to use this extra equipment efficiently?
 - Active Data Guard makes 60+% utilization possible

Possible Hardware Changes

Before FSFO



After FSFO



How to track data state of the standby

- Comparing CURRENT_SCN from v\$database between primary and standby along with SYSTIMESTAMP
- Only accurate as your systems clock drift (ntpd)
- Example

Primary SCN	Primary Time	Standby SCN	Standby Time
4248440	1,193,156,891,753	-	-
4248446	1,193,156,892,764	-	-
4248452	1,193,156,893,774	-	-
4248459	1,193,156,894,786	4248453	1,193,156,894,788

Clients view of Active Data Guard

- How do we inform the client about how far behind the standby is?
- Many possible methods
 - Interceptor layer that rejects queries that exceed window
 - Broadcast current state to the clients
 - Take database out of service when gap exceeds threshold

Our Experience

- It works!!

- Throughput

- Can push hundreds of Megabytes of redo per minute

- Data gap

- Average is less than 1 second (test measurement granularity)
 - Spikes of less than 10 seconds