12TB Data Guard Standby on a Wide Area Network

Kim Khuu

Sr. Database Engineer

Fidelity National Financial

Agenda

- Fidelity National Financial
- HA/DR Requirements
- DR Technical Challenges
- Data Guard Configuration
- Tips & Tricks
    - Creating a 12TB Standby
    - Using the Standby for QA and HotFix testing – open read/write
    - Quickly restoring the original production database after a failover
- Conclusion

248 on the Fortune 500 - Four Major Operating Groups

FIS»
FIDELITY NATIONAL
INFORMATION SERVICES

- Provider of information processing management, outsourcing services, consulting services and application software to the mortgage and financial services industries.

FIDELITY
NATIONAL TITLE GROUP

- Title insurance and escrow services – issues nearly one in three residential and commercial title insurance policies in the U.S.

FIDELITY
NATIONAL PROPERTY AND
CASUALTY INSURANCE GROUP

- Personal insurance products, including residential property, auto, flood and umbrella

Sedgwick CMS

- Claims administration, managed care, program management and cost management services for workers' compensation; disability, FMLA and other employee absence; and liability claims programs

FNF HA/DR Requirements

- Application: **Central clearing house for electronic title and escrow documents**
- High Availability, Performance, & Scalability
  - Addressed by Oracle Real Application Cluster
  - Attend **S281258 –** to get in-depth on how FNF implement Oracle RAC/ASM 10gR2 on Hitachi, AIX Platform
- Disaster Recovery
  - Addressed by Data Guard
    - Recovery Point Objective < 2 hours
    - Recovery Time Objective < 4 hours
    - Zero perceived impact on production database performance

Disaster Recovery Technical Challenges

- Less than 4hrs RTO for a 12TB database composed of over 2000 datafiles.
- High volume of mass transactions
  - Month-end processing generates 1,200 (350GB) archived logs
  - Nightly import generates 60GB archived logs
  - Quarterly archival process, changing inactive files from R/W to R/O
- Additional standby database requirements
  - The need to build a production-like environment for QA stress testing. It’s too expensive to acquire additional hardware.
  - The need to re-instate old primary as standby after a failover

Data Guard Configuration

Primary Database

Oracle 10.2.0.2 RAC, ASM

IBM AIX p670 Servers
4 cpu/16GB memory

Hitachi USP1100 SAN

Standby Database

Oracle 10.2.0.2 RAC, ASM

IBM AIX p570 Servers
3 cpu/12GB memory

Hitachi USP1100 SAN

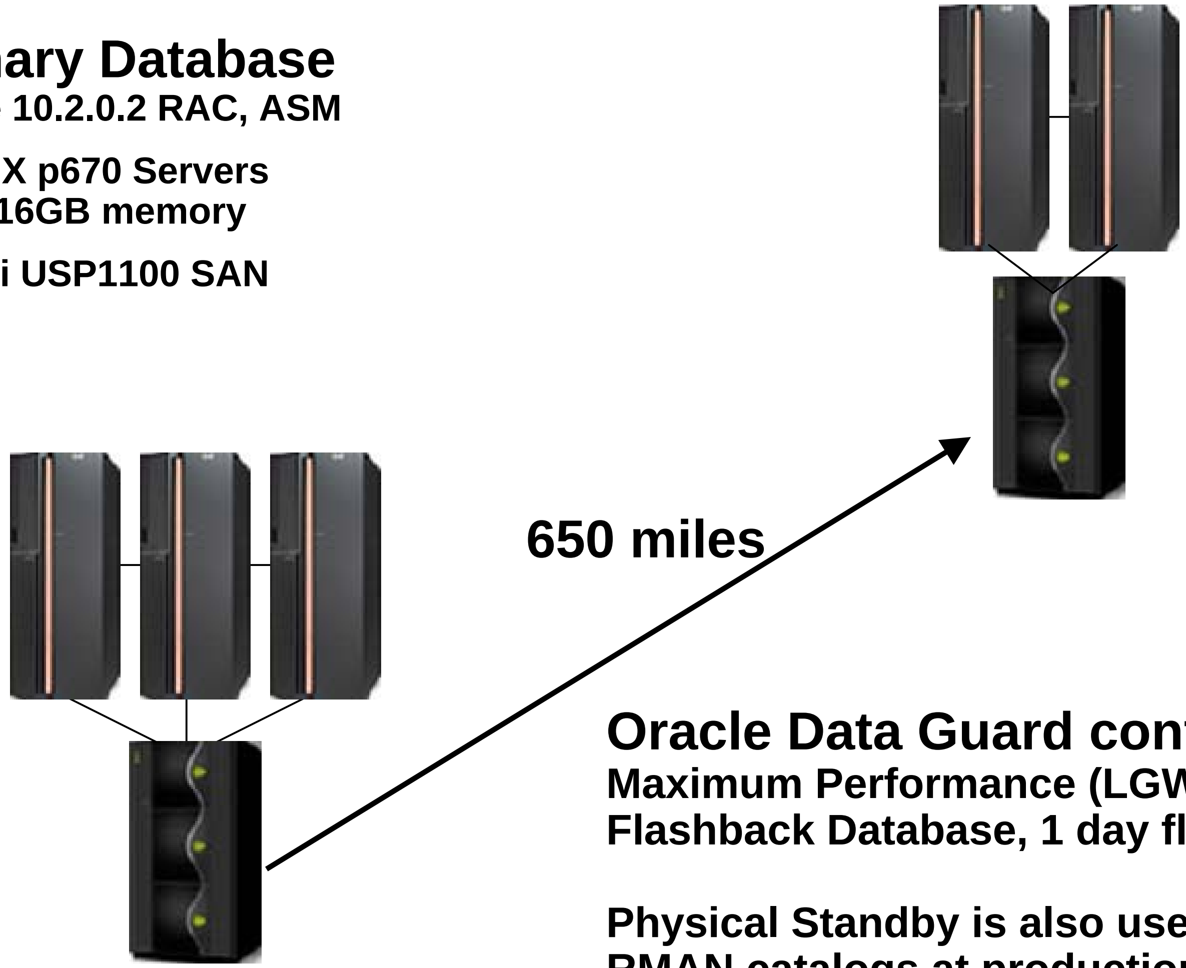

Oracle Data Guard configuration:

Maximum Performance (LGWR ASYNC)
Flashback Database, 1 day flashback retention

Physical Standby is also used to synchronize RMAN catalogs at production and standby sites

Data Guard-Things To Know

- Keep the database in **FORCE LOGGING** mode.
- Consider setting the TIME ZONE for the primary and remote standby databases…if application is sensitive to time.
- If you plan to use ASM, OMF then all databases in Data Guard Configuration should use ASM, OMF consistently.
- Consider add more standby redo logs and not multiplex
- Primary and standby DB_UNIQUE_NAME must be unique in a DataGuard Configuration
- ASM/OMF use the **DB_UNIQUE_NAME** to name the directory and it's part of the file name. Hence must set DB_FILE_NAME_CONVERT and LOG_FILE_NAME_CONVERT
- Be sure to use DB_UNIQUE_NAME to register database to CRS

Data Guard Primary Database Init Parameters

```
DB_NAME                   EDOCPRD
DB_UNIQUE_NAME            EDOCPRD
SERVICE_NAMES             EDOCPRD

LOG_ARCHIVE_CONFIG        ‘DG_CONFIG=(EDOCPRD, EDOCLTC)’
LOG_ARCHIVE_DEST_STATE_1   enable
LOG_ARCHIVE_DEST_1        ‘LOCATION=+DG_EDOC_PF101
                          VALID_FOR=(ALL_LOGFILES, ALLROLES)
                          DB_UNIQUE_NAME=EDOCPRD’
LOG_ARCHIVE_DEST_STATE_2   enable
LOG_ARCHIVE_DEST_2        ‘SERVICE=EDOCPRD_LTC lgwr async
                          VALID_FOR=(ONLINE_LOGFILES, PRIMARY_ROLE)
                          DB_UNIQUE_NAME=EDOCLTC’

FAL_CLIENT                EDOCPRD_CTC
FAL_SERVER                EDOCPRD_LTC

DB_CREATE_FILE_DEST       +DG_EDOC_PD101
DB_FILE_NAME_CONVERT      '+DG_EDOC_PD101/EDOCLTC','+DG_EDOC_PD101/EDOCPRD',
                          '+DG_EDOC_PD501/EDOCLTC','+DG_EDOC_PD501/EDOCPRD'
LOG_FILE_NAME_CONVERT     '+DG_EDOC_PF501/EDOCLTC','+DG_EDOC_PF101/EDOCPRD'

STANDBY_FILE_MANAGEMENT   AUTO
```

Note: Table 13-1 of **Oracle Data Guard Concepts and Administration 10gR2**

Data Guard Standby Database Init Parameters

```
DB_NAME                    EDOCPRD
DB_UNIQUE_NAME             EDOCLTC
SERVICE_NAMES              EDOCLTC

LOG_ARCHIVE_CONFIG         ‘DG_CONFIG=(EDOCLTC,EDOCPRD)’
LOG_ARCHIVE_DEST_STATE_1    enable
LOG_ARCHIVE_DEST_1         ‘LOCATION=+DG_EDOC_PF501
                           VALID_FOR=(ALL_LOGFILES, ALLROLES)
                           DB_UNIQUE_NAME=EDOCLTC’
LOG_ARCHIVE_DEST_STATE_2    enable
LOG_ARCHIVE_DEST_2         ‘SERVICE=EDOCPRD_CTC lgwr async
                           VALID_FOR=(ONLINE_LOGFILES, PRIMARY_ROLE)
                           DB_UNIQUE_NAME=EDOCPRD’


FAL_CLIENT                 EDOCPRD_LTC
FAL_SERVER                 EDOCPRD_CTC

DB_CREATE_FILE_DEST        +DG_EDOC_PD101
DB_FILE_NAME_CONVERT       '+DG_EDOC_PD101/EDOCPRD','+DG_EDOC_PD101/EDOCLTC',
                           '+DG_EDOC_PD501/EDOCPRD','+DG_EDOC_PD501/EDOCLTC'
LOG_FILE_NAME_CONVERT      '+DG_EDOC_PF101/EDOCPRD','+DG_EDOC_PF501/EDOCLTC'

STANDBY_FILE_MANAGEMENT    AUTO
```

Data Guard Processes

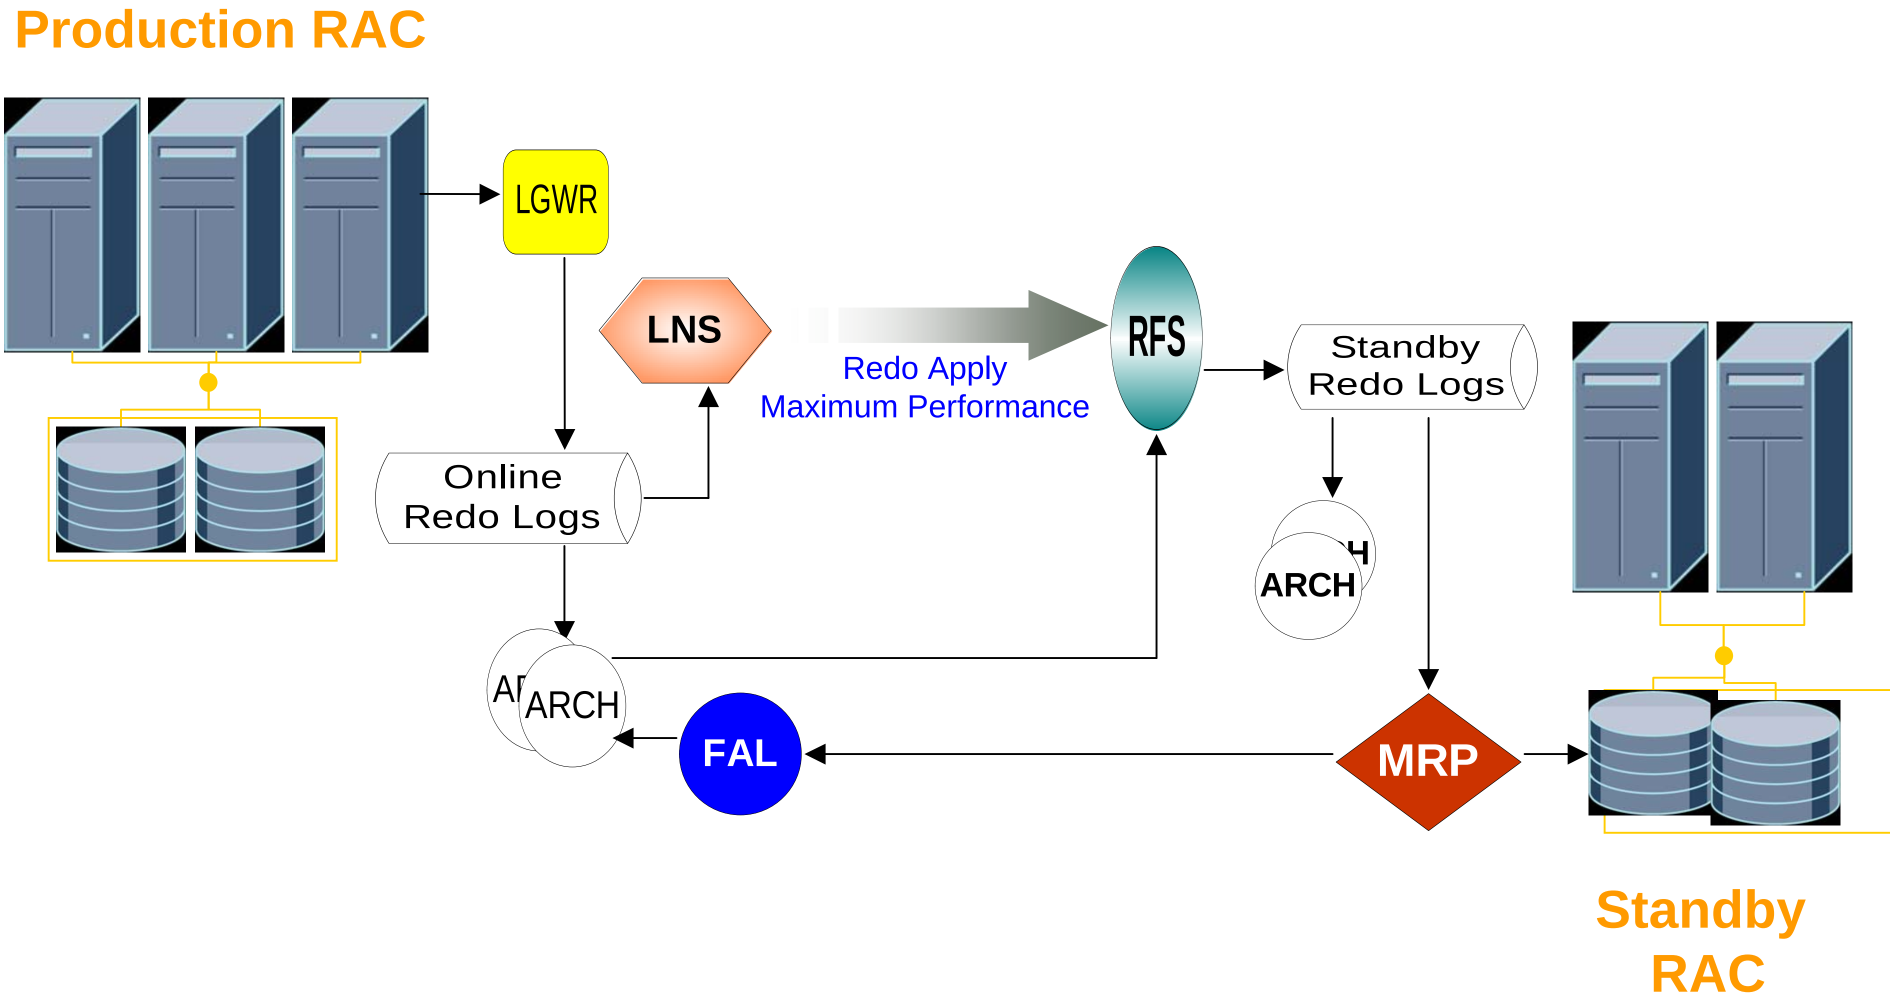

HOW DID WE DO IT?

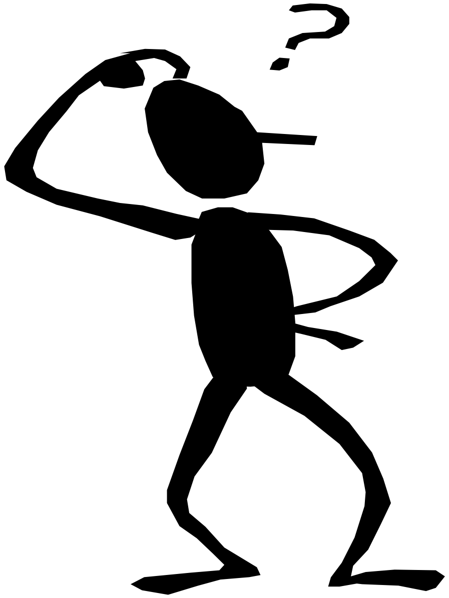

Tips & Tricks

- Creating a 12TB Standby
- Using the Standby for QA and HotFix testing – open read/write
- Quickly restoring the original production database after a failover

Tips & Tricks #1
Creating a 12TB Standby DB

- Design the database with RW and RO partitions
- Backup RO tablespaces only once:
  - backup as compressed backupset incremental level 0 diskratio=0 tag EDOC_backup_Tz_17_JUL_2006 filesperset 5 tablespace PRELIM_XML_2006_Q2 ;
- Use RMAN Level 0 and Level 1 Compressed skip RO to backup database:
  - backup as compressed backupset incremental level 0 skip readonly diskratio=0 tag EDOCPRD1_backup_0z_full_101406 filesperset 5 (database) ;
- Restore the database first
  - RMAN restore database/switch datafile all;
- Restore RO tablespaces ahead of time and RMAN catalog/switch afterward
- In a RAC environment, shutdown other ASM instances during restore
- Do not restore mutilple RO tablespaces concurrently
- Do allocate multiple RMAN channels

Lessons Learned

- Use catalog to backup and restore RO tablespaces as controlfile_record_keep_time is limited
- Consider bigger file size to minimize number of files therefore reduce number of files to backup and restore.
- Can’t specify OMF file name when you have to create data file; use ASM DG name
- Consider Media Server and NetBackup Advanced Client to backup directly to tape with parallel streams

Tips & Tricks #2

Using Standby Database to Stress Test Application and Hot Fixes

Ensure flashback database enabled on standby before open for RW

```
SQL> alter system set db_recovery_file_dest_size = 32g ;
SQL> alter system set db_recovery_file_dest='+DG_EDOC_PF501';
SQL> alter system set db_flashback_retention_target=1440 ;
```

Create Restore Point on standby database

```
SQL> alter database recover managed standby database cancel;

SQL> create restore point before_lt guarantee flashback database;
```

Defer Redo Transport on primary database (all nodes)

```
SQL> alter system set log_archive_dest_state_2=defer;

SQL> alter system archive log current;
```

`Note:  For more information about "before_lt", query v$restore_point`

Using Standby to Stress Test Application & Hot Fixes Part 2 - Open Read/Write

Activate the standby database – open read/write

```
SQL> alter database activate standby database;

SQL> select * from v$database

DATABASE_ROLE    DB_UNIQUE  OPEN_MODE  PROTECTION_MODE  SWITCHOVER_STATUS
---------------- ---------- ---------- ---------------- ------------------
PRIMARY          EDOCTLR    MOUNTED    UNPROTECTED      SESSIONS ACTIVE

SQL> startup mount force;

SQL> alter system set log_archive_dest_state_2=defer;

SQL> alter database open;
```

Using Standby to Stress Test Application & Hot Fixes Part 3 - Revert to Standby

Revert the database back to a synchronized physical standby

```
SQL> startup mount force;
SQL> flashback database to restore point before_lt;
SQL> alter database convert to physical standby;
SQL> startup mount force;

SQL> alter database recover managed standby database disconnect;
SQL> select * from v$database;
DATABASE_ROLE      DB_UNIQUE_  OPEN_MODE   PROTECTION_MODE
-----------------  ----------  ---------   -------------------
PHYSICAL STANDBY   EDOCTLR     MOUNTED     MAXIMUM PERFORMANCE
```

Re-enable Data Guard Transport, on primary (all nodes)

```
SQL> alter system set log_archive_dest_state_2=enable;
```

Note: During standby database open for RW, log shipping are deferred therefore you can have a big GAP. If so, consider to apply an incremental backup created from the primary database.

Lessons Learned

- Consider using an *archive log repository* for 10g databases
  - Can be used to maintain protection while standby is open R/W
- Nice to have:
  - A Data Guard enhancement that would enable redo transport while standby is open R/W

Tips & Tricks #3
Reinstate Old Production DB After Failed Over

After Failover - reinstate the failed primary database as a standby for the new production database using Flashback Database

On the new primary

```
SQL> select to_char(standby_became_primary_scn)from v$database;
```

On the failed primary

```
SQL> shutdown immediate;

SQL> startup mount;

SQL> flashback database to scn <standby_became_primary_scn>;

SQL> alter database convert to physical standby;

SQL> shutdown immediate;

SQL> startup mount;

SQL> alter database recover managed standby database using current
     logfile disconnect;
```

Lessons Learned

- Big Win using Flashback Database to quickly reinstate old primary
  - Avoids 92 hours restoring 1600 backupsets from tapes to disks
  - 115 hours restoring 40 read-only tablespaces and level 0 backupsets
  - Saves a minimum of 8.5 days of standby database outage

Conclusion

Bottom Line:

Data Guard 10g with Flashback Database made it possible to meet FNF operational requirements and service level commitments including disaster recovery objective for our 12TB database.

Additional Data Guard Sessions

- *Session 281207:* Next Generation Oracle Database Availability: A Sneak Preview, Tuesday 10:45am, Moscone South 103
- *Session 281212:* Oracle Data Guard: Defining the Next Era in Data Availability and Data Protection, Tuesday 1:45pm Moscone South 305
- *Session 281210:* Oracle Data Guard Tips and Tricks: Direct from Oracle Development, Wednesday 11:30am, Moscone South 104
- *Session 281208:* MAA Best Practices: Building a Highly Available and Disaster-Proof Architecture, Using Data Guard, Oracle RAC, Automatic Storage Management, and Flashback, Thursday 12:30pm, Moscone South 102
- *Session 281209:* MAA Best Practices: Reducing Downtime for Planned Maintenance Operations Using Oracle Database 10g, Thursday 2:00pm, Moscone South 304