

Graceful Switchover and Switchback for Oracle Standby Database

Lawrence To

lto@us.oracle.com

November 19, 1999

Copyright © Oracle Corporation, 1998. All rights reserved.



Please refer to Note 76450 or CR 268500 for the best practice paper. This presentation assumes that you have a firm understanding of Oracle standby databases. If you need an overview of Oracle 8i standby database, please refer to CR # 270332 and # 260261.

This presentation provides an overview of the graceful switchover and switchback techniques described in the best practice paper. The presentation is used to complement the paper.

Agenda

- **Description of Graceful Switchover and Switchback**
- **Advantages and Disadvantages**
- **Prerequisites**
- **Graceful Switchover Steps**
- **Graceful Switchback Steps**
- **Best Practices**

3-2

Copyright © Oracle Corporation, 1999. All rights reserved.



This presentation gives the audience a quick overview of the details found in the best practice paper.

Description

- **Activating Standby Database in 7.3 and higher will result in**
 - **resetlogs operation**
 - **incompatibility with previous database backups and archives**
 - **incompatibility between Production and Standby Databases**

3-3

Copyright © Oracle Corporation, 1999. All rights reserved.



When activating a standby database with the documented procedures using the *alter database activate standby database* command, Oracle does an implicit resetlogs operation. This operation is required when recovering from a standby controlfile, backup controlfile, creating a new controlfile with the resetlogs operation, or when executing point in time recovery. The resetlogs operation helps prevent users from applying redo that is not compatible to the database. However, after a resetlogs operation, all database backups and the previous production database will not be compatible to the database in the standby database. The resetlogs operation is often considered the point of no return. Any previous backups or any data files requiring recovery prior to the resetlogs operation will be unusable.

Description

- **Graceful Switchover and Switchback**
 - **Technique to safely switch production and standby database roles**
 - **Avoids resetlogs operation**
 - **Standby and production databases are still compatible**
 - **Backups are still compatible and relevant**

3-4

Copyright © Oracle Corporation, 1999. All rights reserved.



In some planned events such as simple system maintenance on the primary site or production machine, a failover to the standby site or standby database does not necessarily require a resetlogs operation if a graceful switchover is accomplished successfully. Graceful switchover and switchback allows production and standby databases to switch roles. Basically the standby database will become the production database and the previous production will become the new standby database. The unique advantage of this technique is that both databases will remain compatible. In other words, the archives from the new production database can be applied by the new standby database. The standby database does not need to be rebuilt from backups from the new production database which is the normal documented case. This big advantage is accomplished only if we avoid the resetlogs operation and follow the best practice steps.

Description

- **Officially supported for standby databases versions 7.3 and higher if best practices are tested and followed**
- **Caution**

A user error can result in corrupting the production or standby databases. Test and validate best practice steps found in CR 268500.

3-5

Copyright © Oracle Corporation, 1999. All rights reserved.



Recently development and COE have worked very closely in documenting the best practice steps. If they are followed, tested and validated, we are confident that graceful switchover and switchback can be executed safely in your production environment. Since scripts and some checks and balances for graceful switchover are not inherent in the current standby product, an user error can result in corrupting or damaging the production or standby database. Extreme caution must be taken.

Description

- **Standby Database Architecture consists of**
 - **Primary Site**
 - **Production Database (single instance or OPS)**
 - **Standby Site**
 - **Standby database that is mounted in recovery mode**
 - **Standby Control File**

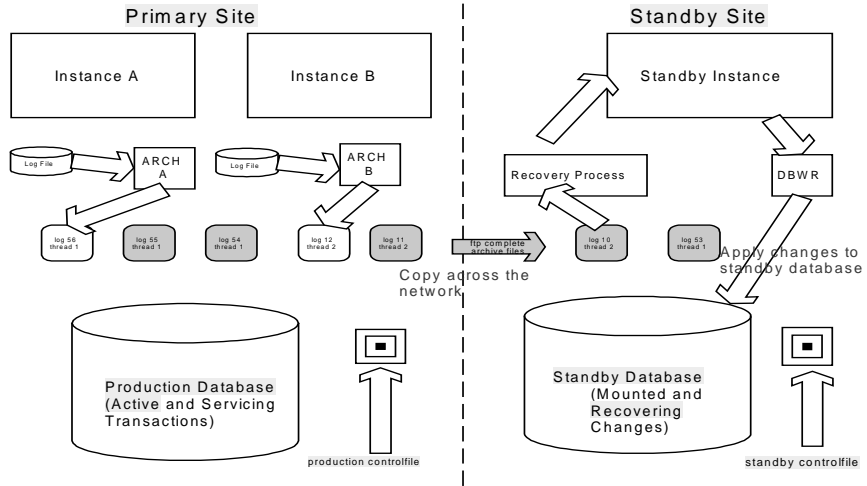
3-6

Copyright © Oracle Corporation, 1999. All rights reserved.



The next two slides give a brief overview of a typical standby architecture.

Standby Database Architecture



3-7

Copyright © Oracle Corporation, 1999. All rights reserved.



Description

- **Graceful Switchover or Switchback**
 - switching production and standby roles
 - ideal for some scheduled outages
 - ideal for unscheduled outages if archive logs and online redo logs are accessible

3-8

Copyright © Oracle Corporation, 1999. All rights reserved.



Advantages

- **Advantages**
 - **Avoids resetlogs operation**
 - **Reduced MTTR in case new production database fails**
 - **No need to recreate new standby database**
 - **Compatibility with previous backups**
 - **Reduced system administration**

3-9

Copyright © Oracle Corporation, 1999. All rights reserved.



Disadvantages

- **User error can result in loss of production or standby databases**
- **Slightly slower fail over times**

3-10

Copyright © Oracle Corporation, 1999. All rights reserved.



The disadvantages described are compared to normal standby database implementations. Typically, opening or activating a standby database will result in an implicit resetlogs operation. The resetlogs will zero out all online redo logs and create a new database version that is not compatible with previous backups and the previous production database.

User error such as copying the online redo logs while the database is still running or applying the online redo logs directly can seriously corrupt the database or prevent the graceful switchover technique from working.

Prerequisites of Graceful Switchover or Switchback

- **Production database is shutdown and not active**
- **All archives are applied to the standby database**
- **Production database's online redo logs are available and accessible**
- **Standby data files are still intact**
- **If relevant, loss of unrecoverable transactions is acceptable**

3-11

Copyright © Oracle Corporation, 1999. All rights reserved.



These are all the prerequisites to allow for graceful switchover. If the prerequisites are not met, you can still activate the standby database with the standard mechanisms described in slide 3.

Prerequisites of Graceful Switchover or Switchback

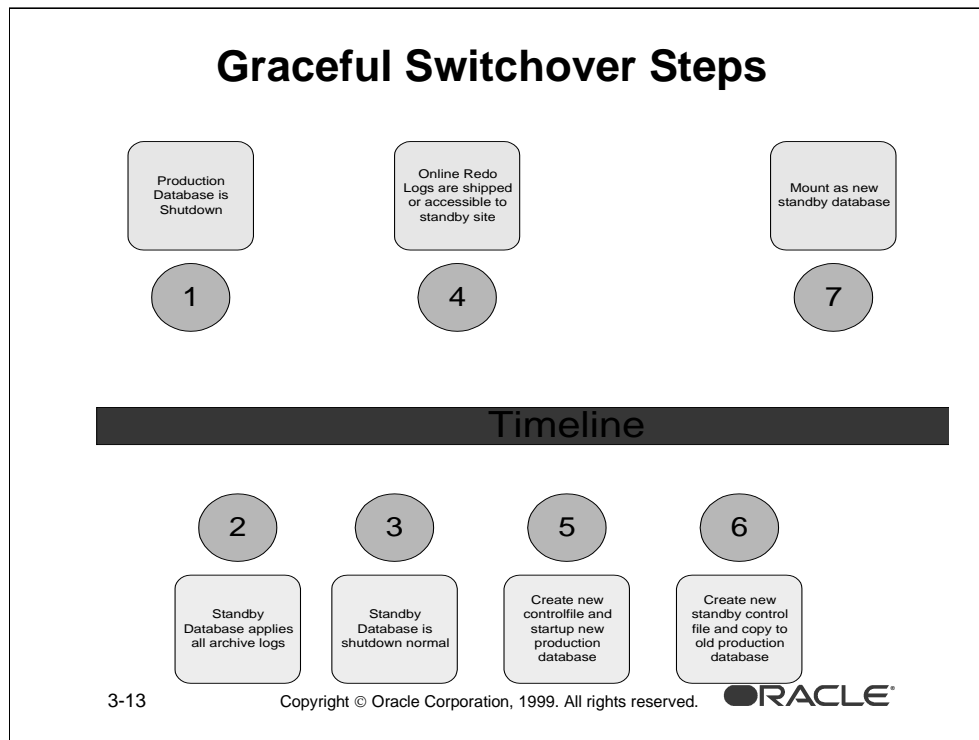
- **All steps have been validated and tested in a test environment first**
- **Experienced DBAs with in-depth knowledge of standby databases are managing the environment**

3-12

Copyright © Oracle Corporation, 1999. All rights reserved.



For an experienced DBA with in-depth recovery knowledge, these steps may seem straightforward; however, for the less skilled individual, these steps may seem confusing and costly errors are more likely. Please ensure that experienced DBAs, consultants or support engineers setup, test, validate and automate these steps to ensure a safe graceful switchover and switchback.



The subsequent slides will describe each step in more detail. In the best practice paper, a table exists that lists a more detailed timeline of steps. Notice that the primary and standby sites are constant. They reflect a location where the production database or standby database initially resided.

Graceful Switchover Steps

- **Primary Site: Step 1 - Shutdown Production Database**
 - **Recommended operation: shutdown normal or shutdown immediate**
 - **If all instances fail, ensure that online redo logs are intact and no media corruption exists.**
 - **Check for critical Oracle errors in production's alert.log (e.g. ORA-600)**

3-14

Copyright © Oracle Corporation, 1999. All rights reserved.



If ORA-600 errors appear in the alert.log, the integrity of the online redo logs may be suspect. Attempting a graceful switchover may in fact corrupt the fail over database. In this case, call support or abort the graceful switchover. Incomplete recovery may be required to safeguard from applying a corrupted archive or online redo log.

If a clean shutdown is not possible, ensure that the instance and all background processes are down. Check for system errors for the online redo logs and archives. If the online redo logs or unapplied archives are damaged, graceful switchover is not possible.

Graceful Switchover Steps

- **Standby Site: Step 2 - Standby Database Applies Archive Logs**
 - Standby database must apply all archive logs successfully
 - Production's online redo logs should **NOT** be applied

```
Startup pfile=initstandby.ora nomount
alter database mount standby database
recover standby database
```

3-15

Copyright © Oracle Corporation, 1999. All rights reserved.



The standby database should apply the archives through normal documented procedures for a standby database.

Graceful Switchover Steps

- **Standby Site: Step 3: Shutdown standby database**
 - **Must complete applying all archives first**
 - **Shutdown normal**



Shutdown

3-16

Copyright © Oracle Corporation, 1999. All rights reserved.

 ORACLE®

After applying all the archives, cancel recovery and shutdown the standby instance.

Graceful Switchover Steps

- **Primary Site: Step 4 - Online Redo Logs are copied to Standby site**
 - **Production database must be shutdown. Online redo logs need to be intact.**
 - **Copy the online redo logs to standby site.**

3-17

Copyright © Oracle Corporation, 1999. All rights reserved.



To prevent the case where you cannot access the online redo logs in case of a primary site failure, customers have remote mirrored the online redo logs. In the case of remote mirrors, the online redo logs are always available to the standby site. One supported technique is using synchronous EMC SRDF to remote mirror the online redo logs from one EMC Symmetrix to a second EMC Symmetrix. The standby machine will access the second Symmetrix machine.

Graceful Switchover Steps

- **Standby Site: Step 5 - Create new controlfile and startup new production**
 - **Create controlfile script needs to be created that points to all the data files and online redo logs**
 - **Noresetlogs and archivelog options need to be utilized**
 - **Fools the standby database from believing this is production**

3-18

Copyright © Oracle Corporation, 1999. All rights reserved.



The create controlfile script must be created from issuing an “alter database backup controlfile to trace” from the PRODUCTION database and NOT the STANDBY database. The script will then need to be modified to point to the correct data file and log file names. Please refer to the best practice paper for more details and for OPS considerations.

Graceful Switchover Steps

- **Standby Site: Step 5 - Create new controlfile and startup new production**
 - Recover database will apply online redo logs
 - Open the database as production

Create Controlfile
Recover database
Alter database open

3-19

Copyright © Oracle Corporation, 1999. All rights reserved.



You should start the database in the standby site with a new “production” init.ora file pointing at different control files. You can then issue the following commands:

```
svrmgrl> connect internal
```

```
svrmgrl> @create_controlfile
```

Within the create controlfile script, commands to the recover the database and open database exist. In a high level, these commands exist in the create controlfile script:

1. Create controlfile
2. Offline datafiles
3. Recover database
4. Alter system archive log all
5. Alter database open
6. Alter database rename file ‘MISSING...’ to ‘filename’

Files in normal offline and read only tablespaces are now named.

7. Alter tablespace online

Online read only tablespaces

8. Add any tempfile entries

Graceful Switchover Steps

- **Standby Site: Step 6 - Create new standby controlfile and copy to new standby database**
 - **Production database now resides on standby site**
 - **Create a new standby controlfile from the production database**
 - **Copy to the new standby database which resides on the Primary Site**
 - **Copy archives across**

3-20

Copyright © Oracle Corporation, 1999. All rights reserved.



At this stage, the new production database is up and running at the standby site. You now need to create a standby controlfile for the new standby database at the primary site.

Graceful Switchover Steps

- **Standby Site: Step 6 - Create new standby controlfile and copy to new standby database**
 - **Issue command on the production database**

```
Alter database create standby controlfile  
as 'filename';
```

Graceful Switchover Steps

- **Primary Site: Step 7- Mount as new standby database**
 - **Old production database will switch roles and become the new standby database**
 - **Init.ora file is modified to point to the new standby controlfile**
 - **Possibly rename datafiles**
 - **Start recovery**

3-22

Copyright © Oracle Corporation, 1999. All rights reserved.



We will now convert the database on the primary site as the new standby database. The trick here is starting with a standby controlfile and commencing recovery. If the data files have different names, you may have to rename them manually with an *alter database rename file 'old filename' to 'new filename'* command. You should create a script and automate this procedure.

You will not be able to open the new standby database as READ ONLY unless you apply at least one archive log.

In step 5, a log switch occurs at database open which is part of the BACKUP CONTROLFILE TO TRACE script. Since at that time, the new standby isn't up yet, the log file is not transmitted to the new standby. In step 6 we created a new standby control file, however this standby controlfile will reflect a different datafile's checkpoint SCN than what is found in the new standby database. Manual recovery will be necessary before attempting to open this database READ ONLY or before using managed recovery.

If you using password files, you may need to recreate the password in the new standby database before successfully mounting it.

Graceful Switchover Steps

- **Primary Site: Step 7- Mount as new standby database**

```
Startup pfile=initStandby.ora nomount  
Alter database mount standby database  
Recover standby database
```

3-23

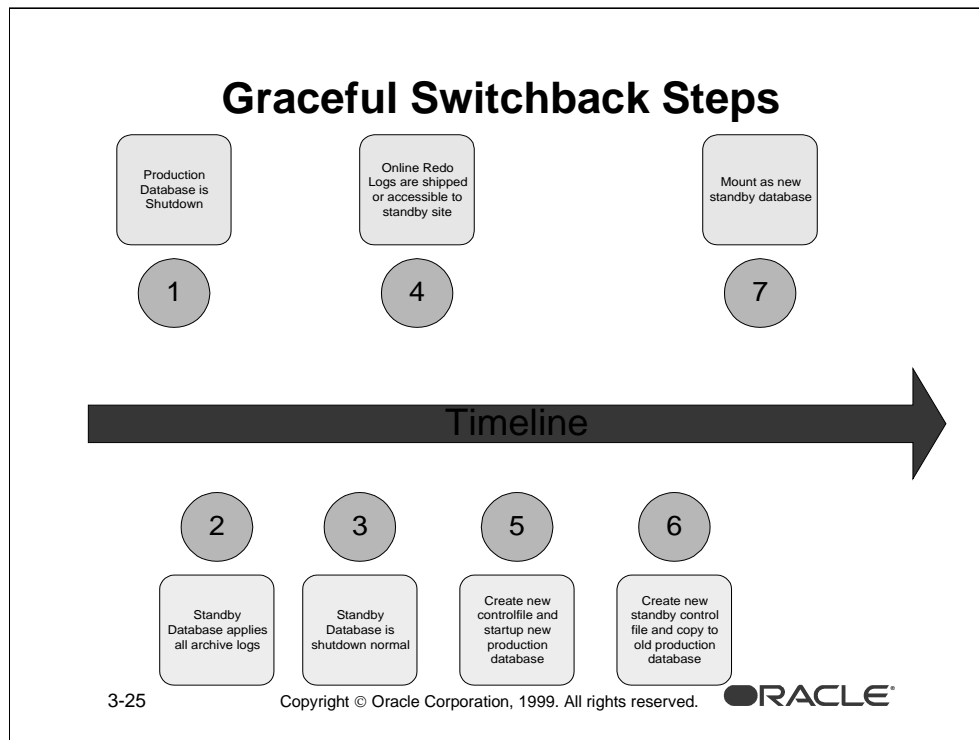
Copyright © Oracle Corporation, 1999. All rights reserved.



After mounting the database, remember you may need to rename the data files if the paths are not the same.

Graceful Switchback Steps

- **Identical to graceful switchover steps**
 - **except target and source databases will be different**



As you can see, the commands are the same. The primary site is now on the bottom and the standby site is on top.

Best Practices

- **Review, validate and test steps found in CR paper 268500**
- **Create a controlled environment that clearly separates standby and production**
- **Automate and script**
- **Rehearse and test scenarios thoroughly**

3-26

Copyright © Oracle Corporation, 1999. All rights reserved.



This presentation only complements the best practice paper. You will still need the best practice paper to successfully implement graceful switchover. The key best practice is rehearse, rehearse, rehearse in a test environment first.

I also suggest building a monitoring environment that detects and reacts to production and standby changes and errors.

Best Practices

- **Monitor system, both databases, and network for errors and activity**
- **Create separate init.ora files for production and standby roles**
- **Validate, modify and check create controlfile script after each database or datafile change**
- **Acquire assistance from COE, Support Services or Consulting**

3-27

Copyright © Oracle Corporation, 1999. All rights reserved.



Good luck and I hope this presentation helps.

I like to thank the following reviewers.

Rick Anderson

Steve Lee