Oracle Maximum
Availability Architecture

Full Stack Role Transition

Oracle ACFS and Oracle Data Guard

ORACLE WHITE PAPER | APRIL 2016

Table of Contents

.

Summary

This Oracle Maximum Availability Architecture (MAA) whitepaper describes how to enable full-stack role transitions by using Oracle Data Guard to replicate Oracle data and the Oracle ASM Cluster File System (ACFS) to replicate non-database files. This MAA solution is best suited for customers and applications requiring:

- The highest level of availability and data protection for the Oracle Database
- The flexibility of a general purpose file system for replicating data that is outside of Oracle
- Full stack switchover for planned maintenance with point-in-time consistency between application and database files
- Full-stack failover for unplanned outage with approximate point-in-time consistency between application and database files

If there is a requirement for exact point-in-time consistency between the database and file system for both planned and unplanned outages, then using Active Data Guard with the Oracle Database File System (DBFS) is the recommended solution. More information on DBFS can be found in the Database SecureFiles and Large Objects Developer's Guide[1].

Data Guard and Active Data Guard[2] provide the management, monitoring, and automation software to create and maintain one or more synchronized copies of a production database to protect Oracle data from failures, disasters, human error, and data corruptions while providing high availability for mission critical applications. Data Guard is included with Oracle Database Enterprise Edition. Active Data Guard is a database option that provides advanced capabilities for data protection and availability.

Oracle ACFS[3] conforms to POSIX standards for Linux and UNIX and to Windows standards for Windows. Oracle ACFS replication supports application files, including executables, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and all other general-purpose application file data. Oracle ACFS includes other advanced features such as file system snapshot, tagging, security, encryption, auditing and highly available NFS (HANFS) services.

This Oracle MAA Best Practices paper is intended for System Administrators and DBAs that have a working knowledge of Oracle ACFS and Oracle Data Guard or Active Data Guard.

[1] http://docs.oracle.com/database/121/ADLOB/adlob_fs.htm#ADLOB45943
[2] http://www.oracle.com/technetwork/database/options/active-data-guard/overview/index.html
[3] http://docs.oracle.com/database/121/OSTMG/GUID-7783FE8B-3BAD-4C4F-83C5-DF1426340290.htm#OSTMG30000

Note that the procedures documented in this paper are identical for both Data Guard and Active Data Guard. All subsequent references to Data Guard are equally applicable to Active Data Guard configurations.

Introduction

The combined use of Data Guard and ACFS with Oracle Database 12c enables the complete replication of production data to a remote copy; both Oracle data and data external to Oracle in an ACFS file system as described in Figure 1. This provides a full-stack solution for transitioning application and database tiers to a replicated copy of production in the event of an unplanned outage (failover) or planned maintenance (switchover).

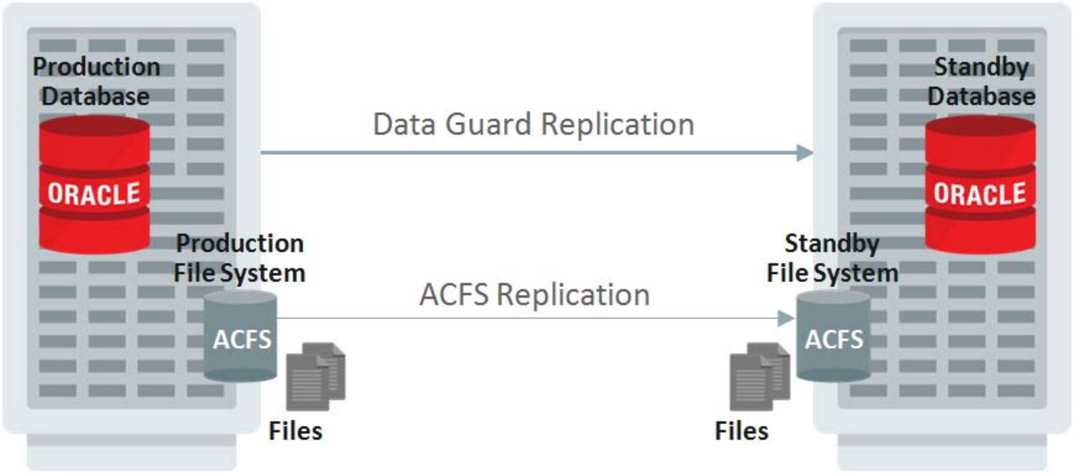

Figure 1: Data Guard and ACFS Replication

Data Guard supports two replication methods.

- Synchronous replication is used by Data Guard Maximum Protection and Maximum Availability protection levels to achieve zero data loss during unplanned outages.
- Asynchronous replication is used by Data Guard Maximum Performance protection level to achieve near-zero data loss protection during unplanned outages.

ACFS only supports asynchronous replication of file system changes. This means that regardless of the Data Guard transport method used to protect the database; unplanned outages always result in approximate point-in-time consistency between database and file system data. This paper describes a procedure to quickly determine the actual degree of consistency between the database and non-database files after an unplanned outage and how to quickly recover to a closer point in time if desired.

Unlike a failover, switchovers for planned maintenance always provide exact point-in-time consistency between database and file system data when using the procedures outlined in this paper. The method of replication, either synchronous or asynchronous, is irrelevant during a planned role transition. All committed transactions and file system changes are replicated to the synchronized copy when using the procedures provided in this paper.

Use Cases

The following use cases involving planned and unplanned outages are documented in this best practice paper:

- Full Stack Role Transition - Planned Maintenance

  Stop the application to perform planned maintenance at the standby system. Execute a Data Guard and ACFS switchover with zero data loss and full consistency between database and file system data. Make the new primary file system read-write and reverse ACFS replication.
- Full Stack Role Transition - Unplanned Outages

  Execute a Data Guard failover and terminate ACFS replication. Compare the point in time that the standby database became primary with the last sync time of the standby and the primary ACFS file system. If the point-in-time consistency of the database and file system are within an acceptable range, then resume service. If the database and file system are at points-in-time that have diverged beyond an acceptable threshold, then follow the procedures documented in this paper to bring them back into an acceptable range and resume service. In both cases, reinstate the old primary as a standby database and reinitialize ACFS replication once it is repaired.

Prerequisites

The use-cases documented in this paper have been validated according to the following prerequisites:

- Oracle Database 11g (11.2.0.4) and Oracle Database 12c (12.1.0.2) with Data Guard.
- Oracle ACFS/Grid Infrastructure 12c (12.1.0.2)
- ACFS file system replication between the primary and standby sites.
- Data Guard broker is highly recommended and is used in the examples that follow.
- Database forced logging is enabled as per standard Data Guard best practice. The following query must return ‘YES’ on both primary and standby databases.

  ```
  SQL> select force_logging from v$database;
  FORCE_LOGGING
  ------------------
  YES
  ```
- Flashback Database

Restrictions and Limitations

Oracle 12c ASM Cluster File System (ACFS) Restrictions and Limitations

- No files associated with the management of Oracle ASM, such as files in the Oracle Grid Infrastructure home and in the Oracle ASM diagnostic directory.
- No support for Oracle Cluster Registry (OCR) and voting files.
- After an ACFS file system has been updated to 12.1.0.2, it can't be reverted to an earlier version nor can it be mounted by an earlier Grid Infrastructure version.
- No support by ACFS mount registry for Oracle Restart configurations[4] .
- Oracle ACFS supports 1023 snapshots on 64-bit systems can be any combination of read-only or read-write.
- Oracle ACFS security does not provide any protection for data sent on the network. If data transmission needs to be secured, then the Oracle Advanced Security option provides enhanced security and authentication.
- Using encryption with database files on Oracle ACFS is not supported.

[4] http://docs.oracle.com/database/121/OSTMG/GUID-FD24F249-E9BB-4E63-B992-11A9868F5D7B.htm#OSTMG94168

- Oracle ACFS supports 64 mounted file systems on 32-bit systems, and 256 mounts on 64-bit systems. However, more file systems can be mounted if there is adequate memory.

The following limits are pertinent to ACFS replication.

- Oracle data files, control files, log files, archives and backups are not supported.
- No support for Oracle Restart.
- Supports only one standby file system for each primary file system.
- Supports eight or fewer nodes mounting the primary file system.
- ACFS standby file system must be empty before initializing the file system.
- Snapshots are not replicated with ACFS replication; however a second set of snapshots can be created on the standby file system to synchronize with the database during unplanned outage.
- When needing to synchronize a file system that is ahead of the database, please note that the file system can be restored to the nearest snapshot available or to a best effort.

Data Guard Configuration and Deployment

The Data Guard configuration in this paper was created using MAA best practices[5] and consists of a 2-node RAC primary and 2-node RAC Standby system. Data Guard is configured using asynchronous transport with Maximum Performance protection mode. Flashback is enabled on primary and standby databases. The Data Guard broker is used to manage the configuration.

```
DGMGRL> show configuration;

Configuration - dg_config

Protection Mode: MaxPerformance
Members:
Chicago - Primary database
Boston - Physical standby database

Fast-Start Failover: DISABLED

Configuration Status:
SUCCESS (status updated 9 seconds ago)
```

ACFS Configuration and Deployment

Oracle ACFS Replication records file updates and changes in real-time and transfers changes asynchronously to the standby file system. The replicated file system at the standby system is available READ ONLY. All Application and meta-data files are written to the primary ACFS file system open read-write.

The primary and standby file system must be running the same operating system and must have similar configuration. ACFS supports configurations where the primary and standby file system may configured differently such as 4 node RAC primary and 2 node RAC standby. MAA best practices, however, recommend that primary and standby sites be configured identically for disaster recovery to achieve the same performance SLAs post role transition.

[5] http://www.oracle.com/technetwork/database/features/availability/oracle-database-maa-best-practices-155386.html

ACFS Replication is setup using instructions provided in the Automatic Storage Management Administrators Guide[6] Since ACFS replication utilizes Oracle Net for transmitting replication logs between primary and standby nodes, the tnsnames.ora must be updated on each node of the cluster as shown in the example below. For higher availability single client access name (SCAN) VIP, can be used in host definition. In addition set the REMOTE_LISTENER initialization parameter in the Oracle ASM instance before initializing replication.

In the $GRID_HOME/network/admin/tnsnames.ora of both systems:

```
ASM_SUN =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST =scan-sun)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = asm_rep)
    )
  )
ASM_MOON =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = -scan-moon)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = asm_rep)
    )
  )
```

On any node of the primary ASM Instance execute the following:

```
SQL> ALTER SYSTEM SET remote_listener='scan-sun:1521’ sid='*' scope=both;
```

On any node of the standby ASM instance execute the following:

```
SQL> ALTER SYSTEM SET remote_listener='scan-moon:1521’ sid='*' scope=both;
```

Add service to ASM Instances on the primary and standby sites:

```
SQL> show parameter service

NAME                              TYPE       VALUE
------------------------------------ ----------- ------------------------------
service_names                     string     +ASM


SQL> alter system set service_names='+ASM,asm_rep' scope=both;


SQL> show parameter service

NAME                              TYPE       VALUE
------------------------------------ ----------- ------------------------------
service_names                     string     +ASM, asm_rep
```

[6] http://docs.oracle.com/database/121/OSTMG/GUID-13F8260D-7549-475E-A701-A12386FE7048.htm#OSTMG94793

All of the examples illustrated in this document use the naming in Table 1:

TABLE 1: NAMING

	Primary	Standby
Hosts	exa01, exa02	exa03, exa04
Database Unique Name	Chicago	Boston
Instance Names	Chicago1, Chicago2	Boston1, Boston2

Full Stack Role Transition with Database and Application Tiers

The high level process for full stack role transition during planned maintenance is:

- Data Guard switchover
- Terminate ACFS replication and make the standby file system read write.
- Reverse ACFS replication.

Prerequisite: All application, mid-tier and activity on the Primary database and file system must cease before the switchover operation is executed.

Step 1: Sync and validate ACFS replication status

```
[root@exa01 appdata]# acfsutil repl sync /appdata
```

```
[root@exa01 appdata]# acfsutil repl info -c /appdata
```

```
Site:                                Primary
Lag Time:                            00:00:00
Primary status:                      Online
Background Resources:                3 of 3 Online
Primary mount point:                 /appdata
Primary Oracle Net service name:     asm_rep
Standby mount point:                 /appdata
Standby Oracle Net service name:     asm_rep
Standby Oracle Net alias:            admin/****@asm_moon
Replicated tags:
Log compression:                     Off
Debug log level:                     2
```

Step 2: On the primary file system make the ACFS file system read-only

Make sure you run the following commands on all RAC nodes in the cluster that currently have the ACFS file system mounted in read write mode.

Check if the file system is currently being used by issuing the following command:

```
[root@exa01]# lsof /appdata
```

If processes are listed using the file system, kill the process using:

```
[root@exa01]# kill -9 ‘pid’
```

Note the current file permissions for the ACFS mount point (this will be needed on page 12 when reversing the ACFS replication direction):

```
[root@exa01]# stat --format '%a' /appdata
```

When there is no activity on the mount point, issue the following command to set the directory permissions to read only mode:

```
[root@exa01]# chmod 700 /appdata
```

Step 3: Validate the Data Guard configuration

```
DGMGRL> show configuration;

Configuration - dg_config
  Protection Mode: MaxPerformance
  Members:
  Chicago - Primary database
  Boston  - (*) Physical standby database

Fast-Start Failover: DISABLED

Configuration Status:
SUCCESS   (status updated 18 seconds ago)
```

Step 4: Perform the Data Guard switchover

```
DGMGRL>SWITCHOVER TO BOSTON;
```

Step 5: Terminate replication and make the standby ACFS file system read-write

In planned maintenance both the primary and standby file systems are available, terminate replication first from primary file system then the standby file system.

```
[root@exa01]# sync
[root@exa01]# acfsutil repl sync /appdata

[root@exa01]# acfsutil repl terminate primary /appdata

[root@exa03]# acfsutil repl terminate standby /appdata

[root@exa03]# acfsutil repl info -c /appdata
Site:                               Standby
Standby status:                     Online
Last sync time with primary:        Mon Oct 12 12:31:36 2015
Background Resources:               4 of 4 Online
Standby mount point:                /appdata
Standby Oracle Net service name:    asm_rep
Primary mount point:                /appdata
Primary Oracle Net service name:    asm_rep
Primary Oracle Net alias:           admin/****@asm_sun
Replicated tags:
```

```
Log compression:                     Off
Debug log level:                     2
```

Step 6: Validate the Data Guard configuration

```
DGMGRL> show configuration;

Configuration - dg_config

  Protection Mode: MaxPerformance
  Members:
  Boston  - Primary database
  Chicago - (*) Physical standby database

Fast-Start Failover: DISABLED

Configuration Status:
SUCCESS   (status updated 37 seconds ago)
```

At this stage the switchover was successful and Boston is the new Primary. Applications can resume processing.

Step 7: Setup ACFS Replication in reverse direction

To reinitialize replication, the standby ACFS file system should be empty and mounted on only ONE node of the cluster.

```
[root@exa01]# chmod XXX /appdata    # Use the permissions recorded above on page 10

[root@exa02]# umount /appdata       # Unmount on all but one RAC node on the standby

[root@exa01]# acfsutil rep init standby -p admin/welcome1@asm_moon /appdata

[root@exa03]# acfsutil rep init primary -s admin/welcome1@asm_sun /appdata
```

Note: Unlike Data Guard switchover, reversing ACFS replication requires copying the entire file system from the new primary site to the old primary site. There is no way to use the original copy of the file system at the original site to resume replication.

Step 8: Mount and validate ACFS file system on other nodes of the standby cluster

```
[root@exa02]# mount -t acfs /dev/asm/appdata-369 /appdata

[root@exa02]# df -h
Filesystem                         Size   Used   Avail  Use% Mounted on
/dev/mapper/VGExaDb-LVDbSys1       30G    7.8G   21G    28% /
/dev/sda1                          496M   40M    431M    9% /boot
/dev/mapper/VGExaDb-LVDbOra1       99G    64G    30G    69% /u01
tmpfs                              95G    628M   94G     1% /dev/shm
/dev/asm/appdata-369               20G    246M   20G     2% /appdata
```

Step9: Validate configuration

```
[root@exa03]# acfsutil repl info -c /appdata
Site:                                Primary
Lag Time:                            00:00:00
```

```
Primary status:                      Online
Background Resources:                3 of 3 Online
Primary mount point:                 /appdata
Primary Oracle Net service name:     asm_rep
Standby mount point:                 /appdata
Standby Oracle Net service name:     asm_rep
Standby Oracle Net alias:            admin/****@asm_sun
Replicated tags:
Log compression:                     Off
Debug log level:                     2


[root@exa01]# acfsutil repl info -c /appdata
Site:                                Standby
Standby status:                      Online
Last sync time with primary:         Mon Oct 12 13:25:48 2015
Background Resources:                4 of 4 Online
Standby mount point:                 /appdata
Standby Oracle Net service name:     asm_rep
Primary mount point:                 /appdata
Primary Oracle Net service name:     asm_rep
Primary Oracle Net alias:            admin/****@asm_moon
Replicated tags:
Log compression:                     Off
Debug log level:                     2
```

Full Stack Role Transition

The high level process for full stack role transition during an unplanned outage is:

- Data Guard Failover
- Terminate ACFS replication and make the standby file system read write.
- Reverse Data Guard and ACFS replication when the original Primary system is available.

Step 1: Execute the Data Guard failover on the standby database

```
DGMGRL>FAILOVER TO BOSTON;
```

Step 2: Terminate ACFS replication on standby

```
[root@exa03]# acfsutil repl terminate standby /appdata
```

At this stage replication is terminated on standby and ACFS file system is Read-Write.

Step 3: Validate the Data Guard configuration

```
DGMGRL> show configuration;

Configuration - dg_config

  Protection Mode: MaxPerformance
  Members:
  boston  - Primary database
    chicago - Physical standby database (disabled)
```

```
      ORA-16661: the standby database needs to be reinstated

Fast-Start Failover: DISABLED

Configuration Status:
SUCCESS   (status updated 59 seconds ago)
```

Step 4: Validate new primary timestamp

On the new Primary determine current timestamp information:

```
SQL> select STANDBY_BECAME_PRIMARY_SCN from v$database;

STANDBY_BECAME_PRIMARY_SCN
--------------------------
                78733893

SQL> select scn_to_timestamp(78733893) from dual;

SCN_TO_TIMESTAMP(78733893)
----------------------------------------------------
16-OCT-15 09.08.28.000000000 AM
```

Step 5: Validate ACFS replication status

```
Site:                                Standby
Standby status:                      Online
Last sync time with primary:         Fri Oct 16 09:08:15 2015
Background Resources:                4 of 4 Online
Standby mount point:                 /appdata
Standby Oracle Net service name:     asm_rep
Primary mount point:                 /appdata
Primary Oracle Net service name:     asm_rep
Primary Oracle Net alias:            admin/****@asm_sun
Replicated tags:
Log compression:                     Off
Debug log level:                     2
```

At this stage the standby ACFS file system is Read-Write, failover was successful and Boston is the new Primary. Applications can resume processing.

Note: It can be seen from the timestamps in this example that the new primary is ahead of the ACFS file system by 13 seconds (09.08.28 - 09:08:15). If this difference was beyond an acceptable threshold for your requirements you may flashback the database to the earlier point in time. Likewise, if the ACFS file system was ahead of the database beyond an acceptable amount you may restore the ACFSS file sytem to an earlier snapshot. Please refer to Appendix A for details of these various techniques if required.

Step 6: Reinstate the old primary database

Once the issues are resolved at the old primary database it can be easily reinstated as a new standby by converting it to a physical standby database (This assumes that the original database files can be recovered and the database restarted):

```
DGMGRL> REINSTATE DATABASE 'CHICAGO’;

DGMGRL>SHOW CONFIGURATION;

Configuration - dg_config
  Protection Mode: MaxPerformance
  Members:
  BOSTON  - Primary database
   CHICAGO - Physical standby database


Fast-Start Failover: DISABLED

Configuration Status:
SUCCESS   (status updated 5 seconds ago)
```

Step 7: Setup ACFS replication in reverse direction

To reinitialize replication, the standby ACFS file system should be empty and mounted on only ONE node of the cluster.

```
[root@exa02]# umount /appdata       # Unmount on all but one RAC node on the standby

[root@exa01]# acfsutil rep init standby -p admin/welcome1@asm_moon /appdata

[root@exa03]# acfsutil rep init primary -s admin/welcome1@asm_sun /appdata
```

Note: Unlike Data Guard reinstantiation, reinitializing ACFS replication copies the entire file system from the new primary site to the old primary site. There is no way to recover the file system at the original site and resume replication.

Step 8: Mount and validate ACFS file system on other nodes of the standby cluster

```
[root@exa02]# mount -t acfs /dev/asm/appdata-369 /appdata

[root@exa02]# df -h
Filesystem                          Size  Used Avail Use% Mounted on
/dev/mapper/VGExaDb-LVDbSys1 30G    7.8G   21G  28% /
/dev/sda1                           496M   40M    431M   9% /boot
/dev/mapper/VGExaDb-LVDbOra1 99G   64G     30G    69% /u01
tmpfs                               95G  628M     94G    1% /dev/shm
/dev/asm/appdata-369                20G  246M     20G    2% /appdata
```

Validate Replication Status:

```
[root@exa03]# acfsutil repl info -c /appdata
Site:                                Primary
Lag Time:                            00:00:00
Primary status:                      Online
Background Resources:                3 of 3 Online
Primary mount point:                 /appdata
Primary Oracle Net service name:     asm_rep
Standby mount point:                 /appdata
Standby Oracle Net service name:     asm_rep
Standby Oracle Net alias:            admin/****@asm_sun
Replicated tags:
Log compression:                     Off
Debug log level:                     2

[root@exa03]# acfsutil repl sync /appdata

[root@exa01]# acfsutil repl info -c /appdata
Site:                                Standby
Standby status:                      Online
Last sync time with primary:         Fri Oct 16 10:25:46 2015
Background Resources:                4 of 4 Online
Standby mount point:                 /appdata
Standby Oracle Net service name:     asm_rep
Primary mount point:                 /appdata
Primary Oracle Net service name:     asm_rep
Primary Oracle Net alias:            admin/****@asm_moon
Replicated tags:
Log compression:                     Off
Debug log level:                     2
```

Operational Best Practices

The following best practices apply when using the combination of ACFS and Data Guard for full-stack failover.

- Use a multi-node RAC configuration to maintain high availability for the failover of application files in the ACFS file system.

- To avoid I/O errors while managing ACFS/ADVM, the volumes must be dismounted before shutting down an ASM instance. The startup sequence must be performed as follows:
  - Start the ASM instance
  - Mount the diskgroups
  - Mount the file system

- If for any reason you can’t unmount the file systems before shutting down ASM, be sure to run the Linux command ‘sync’ twice to flush any cached file system data and metadata to persistent storage.

- It is critical that there is enough disk space available on both sites hosting the primary and the standby file systems to contain the replication logs. To get a good starting point in terms of space requirements execute acfsutil on the primary file system to collect the average rate of change over a 24 hour period with a 15 minute (900 seconds) interval.

  ```
  $ /sbin/acfsutil info fs -s 900 /appdata
  ```

- It is recommended to have at least 24 hours of additional space reserved for both primary and standby file systems in case of an unplanned outage.

- ACFS snapshots are not replicated; hence you cannot use primary file system snapshots to rewind the standby file system. It is recommended that you create standby snapshots at same periodic interval that is setup on the primary in order to have the ability to rewind the standby file system. Additional space is required to create ACFS snapshots.

- The Oracle ACFS parameters AcfsMaxOpenFiles, AcfsMaxCachedFiles control open files and cached files. Normally it’s not required to change the value of AcfsMaxOpenFiles. For example if your parameter is set to 65536 and you are working with 95000 files, you may want to consider increasing the value in your Oracle ACFS file systems or changing AcfsMaxCachedFiles value to get better performance. On Linux x86-64 systems the top limit in defined in /etc/sysctl.conf

  ```
  $ cat /etc/sysctl.conf | grep fs.file-max
  ```

- When terminating replication, first terminate replication on the primary file system, and then the standby file system. If the primary is not available the terminate command can be executed on the standby file system as shown in the unplanned outage example.

- To ensure all changes are sent to the standby file system before terminating the primary site in a planned outage, ensure that all applications are quiesced. Use the Linux sync command then execute the ‘acfsutil repl sync’ command to flush all outstanding replication data from the primary to the standby file system.

  ```
  [root@exa01]# sync
  [root@exa01]# acfsutil repl sync /appdata
  ```

- The Linux Deadline I/O Scheduler is recommended for the disks in the disk group on a Linux system. The Deadline I/O scheduler guarantees a start service time for a request and is enabled by default in Oracle’s Unbreakable Enterprise Kernel. For the Red Hat Compatible Kernel, the default IO scheduler is the 'cfq' scheduler. To verify that Deadline is enabled, issue the following statement as root:

  ```
  [root@exa01]# cat /sys/block/sda/queue/scheduler
  noop [deadline] cfq
  ```

  To set the deadline scheduler issue the following command as root:

  ```
  # echo deadline > /sys/block/hda/queue/scheduler
  ```

Conclusion

Oracle ACFS replication complements Oracle Data Guard as a disaster recovery solution for all non-database files. ACFS replication and Oracle Data Guard provide a comprehensive disaster recovery capability for full stack role transitions addressing both database and application tiers for planned maintenance and unplanned outages.

APPENDIX A – Rewind Techniques

A failover due to an unplanned outage will result in approximate point in time consistency between the new primary database and the ACFS file system. This paper described how to assess the difference in time. If the difference in time is greater than an acceptable threshold, either the database or the file system may be recovered to an earlier point in time to achieve closer alignment. FLASHBACK DATABASE can be used for point-in-time recovery of the Oracle Database to a specific time stamp or SCN. ACFS Snapshots are used to recover the file system to the point in time of a previous snapshot.

After a failover you may compare the point in time that the standby database became primary with the last sync time of the standby ACFS file system. This timestamp is displayed with the acfsutil repl info -c command as ‘Last sync time’ with primary.

Example 1: Database is ahead of file system.

On the new Primary determine current timestamp information:

```
SQL> select STANDBY_BECAME_PRIMARY_SCN from v$database;

STANDBY_BECAME_PRIMARY_SCN
--------------------------
                  78733893

SQL> select scn_to_timestamp(78733893) from dual;

SCN_TO_TIMESTAMP(78733893)
---------------------------------------------------
16-OCT-15 09.08.28.000000000 AM
```

```
On the standby File system gather the last sync time with primary:
[root@exa03 appdata]# acfsutil repl info -c /appdata
Site: Standby
Standby status: Online
Last sync time with primary: Fri Oct 16 09:08:15 2015
Background Resources: 4 of 4 Online
Standby mount point: /appdata
Standby Oracle Net service name: asm_rep
Primary mount point: /appdata
Primary Oracle Net service name: asm_rep
Primary Oracle Net alias: admin/****@asm_primary
Replicated tags:
Log compression: Off
Debug log level: 2
```

As you can see, the Database is ahead of the ACFS file system.

`16-OCT-15 09.08.28` is greater than `16-OCT-15 09:08:15.`

If difference is greater than desired, simply flashback the database to match the time of the file system as determined above:

```
SQL > FLASHBACK DATABASE TO TIMESTAMP TO_TIMESTAMP (‘2015-10-16 09:08:15’,’YYYY-MM-DD
HH24:MI:SS’) ;
```

```
SQL> alter database open resetlogs;
```

Example 2: ACFS file system is ahead of the Database.

On the new Primary determine current timestamp information:

```
SQL> select STANDBY_BECAME_PRIMARY_SCN from v$database;

STANDBY_BECAME_PRIMARY_SCN
--------------------------
                78079874

SQL> select scn_to_timestamp(78079874) from dual;

SCN_TO_TIMESTAMP(78079874)
------------------------------------------------------------
15-OCT-15 11.33.07.000000000 AM
```

Validate ACFS replication status:

```
[root@exa03]# acfsutil repl info -c /appdata
Site:                                Standby
Standby status:                      Online
Last sync time with primary:         Thu Oct 15 11:41:53 2015
Background Resources:                4 of 4 Online
Standby mount point:                 /appdata
Standby Oracle Net service name:     asm_rep
Primary mount point:                 /appdata
Primary Oracle Net service name:     asm_rep
Primary Oracle Net alias:            admin/****@asm_primary
Replicated tags:
Log compression:                     Off
Debug log level:                     2
```

As you can see the ACFS file system is ahead of the Database.

15-OCT-1511:41:53 2015 is greater than 15-OCT-15 11.33.07

Use ACFS snapshot to rewind the file system to a match time from above:

```
[root@exa03]# acfsutil snap info /appdata
snapshot name: 1015_1
snapshot location: /appdata/.ACFS/snaps/1016_1
RO snapshot or RW snapshot: RO
parent name: /appdata
snapshot creation time: Thu Oct 15 10:38:55 2015

snapshot name: 1015_2
snapshot location: /appdata/.ACFS/snaps/1015_2
RO snapshot or RW snapshot: RO
parent name: /appdata
snapshot creation time: Thu Oct 15 11:31:03 2015

number of snapshots: 2
snapshot space usage: 2318336 (2.21 MB)
```

OR using SQL:

```
export ORACLE_HOME=/u01/app/12.1.0.2/grid
export ORACLE_SID=+ASM1
export PATH=$ORACLE_HOME/bin:$PATH

oracle@exa03]$ sqlplus / as sysasm

List out available snapshots using sql below:
   COL fs_name          FORMAT A40
   COL vol_device       FORMAT A30
   COL snap_name        FORMAT A15
   COL create_dtm       FORMAT A11

SELECT
     fs_name
    ,vol_device
    ,snap_name
    ,TO_CHAR(create_time, 'yyyy-mm-dd hh24:mi:ss')
    ,SUBSTR(PARENT,1,10) PARENT
    ,SUBSTR(TYPE,1,4) TYPE
  FROM v$asm_acfssnapshots
```

```
FS_NAME       VOL_DEVICE             SNAP_NAME      TO_CHAR(CREATE_TIME)   PARENT  TYPE

------------ ---------------------- --------------- ------------------- ------- ----

/appdata      /dev/asm/appdata-38   mid_day2       2015-10-09             NULL    RW
                                                   13:57:41

/appdata      /dev/asm/appdata-38   1013_1         2015-10-13             NULL    RO
                                                   12:02:18

/appdata      /dev/asm/appdata-38   1012_2         2015-10-12             NULL    RO
                                                   16:11:22

/appdata      /dev/asm/appdata-38   1012_1         2015-10-12             NULL    RO
                                                   13:44:22
```

Change into the desired snapshot directory and use any utility cp, zip, tar, cpio to copy to the acfs file system.

```
[root@exa01] cd /acfs/.ACFS/snaps/1016_1

[root@exa01] cp /acfs/.ACFS/snaps/1016_1/* /appdata
```

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

- blogs.oracle.com/oracle
- facebook.com/oracle
- twitter.com/oracle
- oracle.com

Integrated Cloud Applications & Platform Services

Full Stack Role Transition - Oracle ACFS and Oracle Data Guad
Author: Ranjit Veen
Contributing Authors: Stephan Haisley, Larry Carpenter