

EMC CLARiiON AX4-5 Storage Replication Consistency Feature with Oracle 11g and ASM Validation Test Report

Copyright © 2008 EMC Corporation. All rights reserved.

Published April, 2008

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

EMC CLARiiON AX4-5 Storage Replication Consistency Feature with Oracle 11g and ASM Validation Test Report

H4356

Introduction	5
iSCSI AX4-5 test environment (Windows).....	5
Host machines.....	5
Storage systems	5
Oracle software.....	5
FC AX4-5 test environment (Linux)	5
Host machines.....	5
Storage systems	6
Oracle software.....	6
Storage system configuration	6
Test scenarios	6
Test details	7
Production host	7
Backup host	8
Test matrix and results	9
Summary	9

Introduction

This report documents functionality testing of the EMC® CLARiiON® AX4-5 SnapView® and MirrorView®/Synchronous replication technology that was done to ensure the AX4-5 readiness in an Oracle Database 11g environment. During replication, data content consistency across multiple LUNs comprising a database must be maintained. As such, testing was focused on using the consistency feature available with SnapView and MirrorView/S in conjunction with Oracle 11g's flashback database and restore point features to create a valid point-in-time backup of an Oracle database that spans multiple LUNs and without having to put the database into hot backup mode during replication.

Both AX4-5 models were tested: iSCSI and Fibre Channel. SnapView is supported on both models and MirrorView/S is supported only on a Fibre Channel model configured with Navisphere® Management Suite. Testing was on Windows 2003 for the iSCSI model and on Linux RHEL 4 for the FC model.

iSCSI AX4-5 test environment (Windows)

Host machines

One Dell PowerEdge 2650 server as the production host with the following components:

Operating system	Microsoft Windows Server 2003 Enterprise Edition Version 5.2 Service Pack 2 Build 3790
RAM	3.75 GB
Naviagent/CLI	6.26.7.0.55

One Dell PowerEdge 6450/700 server as the backup host with the following components:

Operating system	Microsoft Windows Server 2003 Enterprise Edition Version 5.2 Service Pack 1 Build 3790
RAM	3.87 GB
Naviagent/CLI	6.26.7.0.55

Storage systems

One EMC CLARiiON AX4-5i storage array with the following components:

Processors	2
Memory size	1 GB per SP
Number of disks	12 x 400 GB FC @ 10k rpm
Base software	FLARE® 02.23.050.3.526

Oracle software

Oracle Database version	Oracle 11g Release 1
-------------------------	----------------------

FC AX4-5 test environment (Linux)

Host machines

Two Dell 2850 servers, one for production and one for backup host, with the following components:

Operating system	RHEL 4 AS U4 Kernel 2.6.9-34.ELsmp
Memory	8 GB
CPU	2 x Xeon 3.2 GHz

Naviagent/CLI	naviagentcli-6.26.5.0.71-1
PowerPath®	EMCpower.LINUX-5.0.0-156

Storage systems

One EMC CLARiiON AX4-5 storage array for the primary system with the following components:

Processors	2
Memory size	1 GB per SP
Number of disks	12 x 146 GB @ 15k rpm SAS drives
Base software	FLARE 02.23.050.3.530

One EMC CLARiiON AX4-5 storage array (MirrorView/S only) for the secondary system with the following components:

Processors	2
Memory size	1 GB per SP
Number of disks	12 x 146 GB @ 15k rpm SAS drives
Base software	FLARE 02.23.050.3.530

Oracle software

Oracle Database version	Oracle 11g Release 1
ASM	oracleasm-2.6.9-34.ELsmp-2.0.2-1

Storage system configuration

Six 4+1 RAID 5 LUNs were created on the CLARiiON storage array and configured to be managed by Automatic Storage Management (ASM), Oracle's file system and volume manager. These six production LUNs were distributed between four disk groups (using external redundancy) as follows:

DATA_DGRP: Consist of two LUNs for all database files and control files

FLASH_DGRP: Consist of one LUN for the flash recovery area

ARCHIVE_DGRP: Consist of one LUN for the archive logs

REDO_DGRP: Consist of two LUNs for the online redo logs

Snapshots, clones, and remote mirrored pairs (Linux only) were created on the appropriate primary and secondary storage arrays. For MirrorView/S replication, two consistency groups were created out of the six remote mirrored pairs. One consistency group consists of just the mirrored pair for the archive logs and the other consistency group consists of the remaining five mirrored pairs.

Test scenarios

An online backup test that simulates bank transactions of about 400,000 rows being updated over a period of time was used for this purpose. For each test run, updates were ongoing when the replicated copy was split from its source. To further stress test the system, ASM disk group rebalancing was in progress with ongoing updates when the replicated copy was split from its source. The two scenarios tested for each replication technology were:

- Storage consistency, flashback on, database not in hot backup mode, no disk group rebalance
- Storage consistency, flashback on, database not in hot backup mode, with disk group rebalance

To ensure no timing issue, each test scenario was repeated a minimum of five times.

Test details

To enable flashback as well as archiving to a location outside the flash recovery area, the following was added to the init*.ora parameter files of the production and backup instance:

```
DB_RECOVERY_FILE_DEST_SIZE = 100G
DB_RECOVERY_FILE_DEST = '+FLASH_DGRP'
LOG_ARCHIVE_DEST_1 = 'LOCATION=+ARCHIVE_DGRP'
```

Archive logging and flashback were then enabled using the following SQL*Plus commands:

```
ALTER DATABASE ARCHIVELOG;
ALTER DATABASE FLASHBACK ON;
```

It is necessary to archive to a separate location outside the flash recovery area so that the disk group containing the archive logs could be split after the backup was taken. This archived log can then be used to roll a backup forward in time. The following sections details steps executed on the production host to capture a point-in-time copy of the database and on the secondary host to recover and roll the database forward:

Production host

1. While the production database was actively being updated but before taking a snapshot or fracturing the replicated LUNs, create a restore point using the following SQL*Plus command:

```
CREATE RESTORE POINT bkup1 GUARANTEE FLASHBACK DATABASE;
```

A “restore point” is a name (example, *bkup1*) associated with a committed System Change Number (SCN) of the database corresponding to the time the restore point was created. This restore point can later be used to flash the database back to the time of creation of the restore point without the need to determine the SCN.

For tests involving ASM disk group rebalancing, start the rebalance process right after creating the restore point. The following SQL*Plus commands initiate the rebalance process:

```
ALTER DISKGROUP DATA_DGRP REBALANCE POWER 2 NOWAIT;
ALTER DISKGROUP REDO_DGRP REBALANCE POWER 2 NOWAIT;
ALTER DISKGROUP FLASH_DGRP REBALANCE POWER 2 NOWAIT;
ALTER DISKGROUP ARCHIVE_DGRP REBALANCE POWER 2 NOWAIT;
```

2. Once a restore point has been captured, start a snapshot session or fracture of the replicated LUNs. The LUN containing the archive logs is not part of this step. Clones and mirrors were verified to be in consistent or synchronized state before the fracture. The following host-based Navisphere Secure CLI (naviseccli) command creates a replica of the production LUNs for the different replication options:

SnapView snapshot to start a consistent snapshot session

```
naviseccli -User admin_user -Password admin_passwd -scope 0 -h SP_name snapview -startsession session_name -snapshotname lun1,lun2,lun3,lunN -persistence -consistent
```

SnapView clone to start a consistent clone fracture

```
naviseccli -User admin_user -Password admin_passwd -scope 0 -h SP_name snapview -consistentfractureclones -CloneGroupNameCloneId name1 cloneid1 name2 cloneid2 ... nameN cloneidN
```

MirrorView/Synchronous to start a consistent group fracture

```
naviseccli -User admin_user -Password admin_passwd -scope 0 -h SP_name mirror -sync -fracturegroup -name group_name
```

At no point was the production database put in hot backup mode before it was snapped or fractured.

3. After the production database has been snapped or fractured, perform a log switch and capture the highest SCN to be used for recovery:

```
ALTER SYSTEM ARCHIVE LOG CURRENT;
SELECT 'NextChange', next_change# from v$log_history where recid=(select max(recid) from v$log_history);
```

4. After the log switch completes, the LUN containing the archive log is then snapped or fractured depending on the replication option selected. The same Navisphere Secure CLI commands were used but with just the archive log LUN.

Backup host

1. In order for a backup host to access the replicated database, the snapshots, fractured clones, or secondary images must be assigned to a storage group that is connected to the backup host. Use the following method to enable a backup host to access the replicated database:

SnapView snapshots

- Add the snapshots to a storage group that is connected to the backup host.
- Issue applicable operating system commands to scan for new devices, making them visible to the backup host.
- Activate the snapshot to enable the backup host to access the newly visible devices.
`admsnap activate -s session_name`

SnapView clones

- Add the clones to a storage group that is connected to the backup host.
- Issue applicable operating system commands to scan for new devices, making them visible and accessible to the backup host.

MirrorView/S secondary images

Using mirror promote:

- Perform a local consistency group promote of the secondary images that are on the secondary array.
`naviseccli -User admin_user -Password admin_passwd -scope 0 -h SP_name mirror -sync
-promotegroup -name group_name -type local`
- Add the promoted LUNs to a storage group that is connected to the backup host.
- Issue applicable operating system commands to scan for new devices, making them visible and accessible to the backup host.

Using SnapView snapshots:

- Start a consistent snap session.
 - Add the snapshots to a storage group that is connected to the backup host.
 - Issue applicable operating system commands to scan for new devices, making them visible to the backup host.
 - Activate the snapshot to enable the backup host to access the newly visible devices.
2. Mount the disk groups using the following SQL*Plus commands so that they are available to the Oracle database instance running on the backup host:

```
ALTER DISKGROUP DATA_DGRP MOUNT;
ALTER DISKGROUP REDO_DGRP MOUNT;
ALTER DISKGROUP FLASH_DGRP MOUNT;
ALTER DISKGROUP ARCHIVE_DGRP MOUNT;
```

3. Restore and recover the database using the following SQL*Plus commands:

```
SQL> STARTUP MOUNT;
SQL> FLASHBACK DATABASE TO RESTORE POINT bkup1;
SQL> RECOVER AUTOMATIC DATABASE UNTIL CHANGE scn USING BACKUP CONTROLFILE;
```

```

SQL> SHUTDOWN
SQL> STARTUP MOUNT
SQL> ALTER DATABASE OPEN RESETLOGS;
SQL> EXIT

```

Test matrix and results

Feature	Test	Results			Times run
		Start snap session	Clone fracture	MV/S fracture	
CONSISTENT	Flashback ON Database NOT in HOT BACKUP mode	Successfully restored and rolled forward	Successfully restored and rolled forward	Successfully restored and rolled forward	5
	Flashback ON Database NOT in HOT BACKUP mode Disk group rebalance in progress	Successfully restored and rolled forward	Successfully restored and rolled forward	Successfully restored and rolled forward	5

Summary

Test results show that the consistency features of SnapView and MirrorView/S for the CLARiiON AX4-5 storage array, when used in conjunction with Oracle 11g's flashback database and restore point features, enable creation of a usable and valid Oracle database backup. This combination of EMC and Oracle technologies to create a consistent backup eliminates the need to put the database in hot backup mode. Because there is no need to put the database in hot backup mode prior to replication, the production database is not impacted and the Oracle database replication process is simplified.