Leveraging DS8000 Series Advanced Copy Services for Oracle User-Managed Backup and Recovery

Oracle RAC 10g with ASM

Author:
Betty Mason
IBM Systems and Technology Group
Open Systems Lab, San Jose, California
October 2008

Version Oct 14, 2008

Table of contents

Abstract

This white paper explains the use of IBM System Storage DS8000 Advanced Copy Services for manually backing up and recovering Oracle Real Application Clusters 10g (RAC) databases that use Oracle Automatic Storage Management (ASM).

Introduction

This white paper documents the use of IBM® System Storage™ DS8000™ Series Advanced Copy Services for backing up and recovering Oracle Real Application Clusters 10g (RAC) databases using Oracle Automatic Storage Management (ASM). Both Oracle RAC and ASM technologies introduce some differences with regard to Oracle Backup and Recovery when using IBM FlashCopy® technologies, as compared to single-instance, filesystem-based Oracle configurations. Deploying DS8000 Advanced Copy Services can be highly effective in any production shop that runs Oracle RAC 10g databases on ASM. This white paper shows the robustness, speed, and ease with which you can backup, restore and recover databases, no matter their size.

The procedures are based on an Oracle User Group white paper that is titled *Backup and Recovery Generic Guidelines with Point-in-Time Copy Solutions and Automatic Storage Management* (located at: http://ioug.itconvergence.com/pls/apex/CILOUG.download_my_file?p_file=403).

This white paper demonstrates the use of Advanced Copy Services FlashCopy and Metro Mirror to do the following:

- Backup and restore an Oracle RAC database with Oracle Database 10g ASM
- Clone an Oracle RAC database with Oracle ASM to a local Oracle RAC
- Clone an Oracle RAC database with Oracle ASM to a remote DS8000 and Oracle RAC
- Backup and recover the Oracle RAC cluster services information

For demonstration purposes, 4-node Oracle RAC clusters were used. The process was repeated on the Linux and AIX® operating systems to demonstrate the operating-system (OS)-independence of the procedures. On all of the test beds, Oracle Clusterware, ASM and database binaries were installed in separate local directories. On the DS8000, volumes were allocated for the Oracle Cluster Registry, voting disks and the ASM diskgroups. The volumes were created by using a consistent naming convention for understanding the relationships between DS8000 volumes, DS8000 copy relationships and the ASM diskgroups. One DS8000 volume was allocated for each Oracle RAC cluster to serve as the mountable backup disk. Consistent naming allows the user to know at a glance what portion of the database is being acted upon in the backup and recovery scenarios.

Audience

The intended audience is any Technical Lead, System Administrator, Storage Administrator, or Oracle Database Administrator in a production environment who is experienced with backup and recovery concepts and procedures. After reading this white paper, the technical staff will understand how to backup and recover an Oracle RAC 10g database with ASM on DS8000 using Advanced Copy Services. The Backup and Recovery by Example section of this paper is written in a step-by-step manner to explain the

process and, thus, can be integrated into any established backup and recovery strategy. A complete list of references will assist the reader in researching any topic requiring additional information.

Technology

In this section, you will find overviews for the concepts related to DS8000 storage attributes and Oracle RAC components. This text describes database backup and recovery and the Oracle components and commands used for the backup and recovery scenarios.

IBM System Storage DS8000 Series overview

The IBM System Storage DS8000 Series is a high-performance, reliable, and exceptionally scalable disk storage system designed to:

- Deliver robust, flexible, and cost-effective disk storage for the mission-critical workloads of medium and large enterprises.
- Enable creation of multiple Storage System logical partitions (LPARs) in a single DS8300 Model 9A2 that you can use for separate production, test, or other unique storage environments.
- Support high availability, storage sharing, and consolidation for a wide variety of operating systems and mixed server environments.
- Help increase storage administration productivity with centralized and simplified management.

The DS8000 series' Advanced Copy Services functions (FlashCopy, Metro Mirror, Global Copy, and Global Mirror) provide a tool set to design business continuity and Remote solutions.

Backup

FlashCopy does not impact your backup time, but it allows you to create a point-in-time (PiT) copy consistent backup across volumes, with a minimum of interruption for your source host. The FlashCopy target(s) can then be mounted on a different host or the backup server. With this procedure, backup speed is less important, because the backup time does not require service interruption for the host, which is dependent on the source volumes.

Restore

Maintaining FlashCopy relationships enables very fast restore of a PiT consistent backup. When an incremental background copy process has completed, and a restore is needed, it is possible to reverse the FlashCopy mappings to very quickly restore the PiT consistent backup from the preceding FlashCopy.

Moving and migrating data

When you need to move a consistent dataset from one host to another, FlashCopy can facilitate this action with a minimum of interruption for the host application which is dependent on the source volume. The FlashCopy target is immediately available, as is the source volume.

FlashCopy

The FlashCopy feature enables you to create full volume copies of data. When you set up a FlashCopy operation, a relationship is established between source and target volumes, and a bitmap

of the source volume is created. Once this relationship and a bitmap are created, the target volume can be accessed as though all the data had been physically copied. While a relationship between the source and target volume exists, a background process copies the tracks from the source to the target volume. FlashCopy is also known as PiT copy. This technique is used to help solve the problem that it is difficult to make a consistent copy of a volume, which is being constantly updated. When FlashCopy is invoked, the resulting data at the target appears as if the copy were made instantly.

Benefits of FlashCopy

FlashCopy creates a PiT copy that is typically used when you need to create a copy of the production data with little or no application down time (depending on the application). You can use it for online backup, testing new applications, or creating a database for data-mining uses. The copy looks like the original source volume and is an instantly available, binary copy.

Consistency for FlashCopy

After a FlashCopy mapping is started, the source and target volumes are, by definition, consistent. The mapping represents a single point in time and every FlashCopy mapping in a consistency group represents the same point in time.

Remote Mirror and Copy features

The Remote Mirror and Copy feature is a flexible data mirroring technology that allows replication between volumes on two or more disk storage systems. DS8000 storage units can participate in remote mirror and copy solutions with the ESS Model 750, ESS Model 800 and DS6000 storage units.

Metro Mirror

Metro Mirror provides real-time mirroring of logical volumes between two DS8000 subsystems that can be located up to 300 km from each other. It is a synchronous copy solution where the attached hosts' write operations must be completed on both logical volume copies (local and remote site) before they are considered to be complete.

Global Copy

Global Copy copies data asynchronously over longer distances than is possible with Metro Mirror. When operating in this mode, the source volume sends a periodic, incremental copy of updated tracks to the target volume instead of a constant stream of updates. This causes less impact to application writes for source volumes and less demand for bandwidth resources, but allows a more flexible use of the available bandwidth.

Global Mirror

Global Mirror provides a long-distance, remote copy feature across two sites by using asynchronous technology. Global Mirror operations provide the following benefits:

- Support for virtually unlimited distance between local and remote sites, with the distance typically limited only by the network capabilities and the channel extension technology. This unlimited distance enables you to choose your remote site location based on business needs and enables site separation to add protection from localized disasters.

- A consistent and restartable data copy at the remote site, created with minimal impact to applications at the local site.

Oracle RAC components overview

This section provides a summary of the Oracle components that are used for the backup and recovery testing.

Oracle ASMLib for LINUX

The Oracle Automatic Storage Management library (ASMLIB) driver simplifies configuration and management of disk devices by removing the need to rebind raw devices that are used with Oracle ASM whenever the system restarts.

Oracle Automatic Storage Management (ASM)

Oracle ASM is a volume manager and a file system for Oracle database files that supports single-instance Oracle Database and Oracle Real Application Clusters (Oracle RAC) configurations. Oracle ASM is the Oracle recommended storage-management solution that provides an alternative to conventional volume managers, file systems, and raw devices.

Oracle ASM diskgroups

Oracle ASM uses diskgroups to store datafiles. An Oracle ASM diskgroup is a collection of disks that Oracle ASM manages as a unit. Within a diskgroup, Oracle ASM exposes a file-system interface for Oracle database files. The content of files that are stored in a diskgroup are evenly distributed, or striped, to prevent hot spots and to provide uniform performance across the disks.

Oracle Clusterware (required by Oracle RAC)

This section provides a summary of the Oracle Clusterware components required by Oracle RAC.

Oracle Real Application Clusters (RAC): Multiple Instance Systems

Some hardware architectures (that is, shared-disk systems) let multiple computers share access to data, software, or peripheral devices. Oracle RAC takes advantage of such architectures by running multiple instances that share a single physical database. Oracle Clusterware is required to run Oracle RAC.

Oracle Cluster Ready Services (CRS)

CRS is a component of Oracle Clusterware. It primarily manages high-availability operations in a cluster. Whatever the CRS process manages is known as a cluster resource. The CRS process manages cluster resources, based on the resource's configuration information that is stored in the Oracle Cluster Registry (OCR). This includes start, stop, monitor, and failover operations. The CRS process generates events when a resource status changes. CRS monitors the Oracle instance, listener, and so on, and automatically restarts these components when a failure occurs. By default, the CRS process attempts to restart a resource five times and then makes no further attempts.

Event management (EVM)

EVM is a background process that publishes events that CRS creates.

Oracle Notification Service (ONS)

ONS is a publish-and-subscribe service that Oracle Clusterware uses for publishing fast application-notification (FAN) events.

RACG

RACG extends Clusterware to support Oracle-specific requirements and complex resources. It runs server-callout scripts when FAN events occur.

Process Monitor Daemon (OPROCD)

This process is locked in memory to monitor the cluster and to provide I/O fencing. It performs its check, stops running and, if the wake up is beyond the expected time, resets the processor and reboots the node. An OPROCD failure results in Oracle Clusterware restarting the node. OPROCD uses the hangcheck timer on Linux platforms.

Oracle Cluster Synchronization Services (CSS)

CSS is the component of Oracle Clusterware that handles group membership for the cluster. Oracle ASM requires the use of CSS to enable synchronization between an Oracle ASM instance and the database instances that rely on it for database file storage.

Overview of Oracle Database Backup and Recovery

Oracle provides various mechanisms for the following processes:

- Database recovery that is required because of different types of failures
- Flexible recovery operations that suit any situation
- Availability of data during backup and recovery operations so that users of the system can continue to work

Types of database failures

Several circumstances can halt operation of an Oracle database, including those described in Table 1. Media (disk) failure is the subject covered in the test scenarios.

Failure	Description - Definition
User error:	Requires data to be recovered to a point in time before the error occurred.
Statement failure:	Occurs when there is a logical failure in the handling of a statement in an Oracle program. When statement failure occurs, any effects of the statement are automatically undone by Oracle and control is returned to the user.
Process failure:	Results from a failure in a user process accessing Oracle, such as an abnormal disconnection or process termination. The background process PMON automatically detects the failed user process, rolls back the uncommitted transaction of the user process, and releases any resources that the process was using.
Instance failure:	Occurs if a problem arises that prevents an instance from continuing work. It can result from a hardware problem such as a power outage, or a software problem such as an operating system failure. After an instance failure, Oracle automatically performs instance recovery.
Media (disk) failure:	An error can occur when trying to write or read a file on disk that is required to operate the database. For example, a disk failure requires you to restore database files from backup media and then perform media recovery. Media recovery must be initiated by the user.

Table 1. Types of database failures

Oracle provides for complete media recovery from all possible types of hardware failures, including disk failures. Options are provided so that a database can be completely recovered or partially recovered to a specific point in time.

Oracle components and commands for backup and recovery scenarios

The following are definitions of the required components and commands to recover from an instance or disk failure:

Database datafiles: Structures containing the actual data being managed by the database.

Redo Log: The redo log is a set of files that contain changes made to the data within the database. The redo log can consist of the online redo log and the archived redo log.

Archived redo log: The filled online redo log files are archived to the archived redo logs before they are reused. If running in Archive log mode, the database can be backed up while it is open and available for use. However, additional administrative operations are required to maintain the archived redo log.

Control Files: The controlfiles include information about the file structure of the database and the current log sequence number. During normal recovery procedures, the information in a controlfile guides the automatic progression of the recovery operation.

Database Backups: Because one or more files can be physically damaged as the result of a disk failure, media recovery requires the restoration of the damaged files from the most recent operating system backup of a database.

Parameter Files: Parameter files contain a list of configuration parameters for that instance and database. The parameter file can be either a text init.ora file or a binary format server parameter file (spfile).

Hot backup mode: Data files are marked with an indicator that shows recovery is needed from x point forward. Transactions cause entire database blocks to be written in full to the redo logs. Putting a database in hot backup mode allows the use of FlashCopy to capture the PiT copy.

Archive log current: In an Oracle RAC environment, writes current redo logs for all threads to the archived logs.

oracleasm: The *oracleasm* support package provides the utilities used to get the ASM driver up and running.

srvctl: Oracle-supplied tool to control resources registered with Oracle Clusterware. The *srvctl* command should only be executed by the database owner (oracle userid).

ocrconfig: Oracle-supplied command to export, import, or restore backups of the OCR disks containing cluster information about registered services and Oracle RAC cluster nodes.

crs_stat: Oracle supplied utility used to show the status of the cluster resources.

RMAN: Oracle Recovery Manager is an Oracle-supplied utility for backup and recovery. RMAN is the Oracle-preferred method for backup and recovery.

SQL*Plus: Command Line tool used to interface with the instance, database, and all components contained in the database.

asmcmd: This is an Oracle supplied utility used to manage ASM disks.

Configuration and setup

This section contains details regarding the configuration and setup for the scenarios tested.

Lab hardware and software configurations

This section explains the hardware and software configurations used in the lab.

Oracle RAC host nodes

Table 2 describes the three 4-host node clusters used in the Backup and Recovery exercises.

RAC server type	Four IBM System x3455	IBM eServer™ pSeries® p690
Processor	Two Dual-Core AMD Opteron™ processor 2222 S 2992.634 MHz	Four
Memory	4 GB	4 GB
Host bus adapter (HBA) model	QLE2462	Emulex LP9002 HBA
Operating system	Red Hat Enterprise Linux® (RHEL) AS4 U4	IBM AIX® 5.3
Kernel version	2.6.9-42.ELsmp	AIX 5.3 ML6
Multipath software	device-mapper-multipath-0.4.5-16.1.RHEL4	Subsystem device driver - path control module (SDD-PCM)
HBA driver	8.01.06	IBM
HBA firmware	4.00.23	02E01974
Cluster file system	None	None
Oracle software	Oracle Database, ASM and Clusterware 10.2.0.3	Oracle Database, ASM and Clusterware 10.2.0.3
Logical configure	4-node RAC members ALPHA, BETA, GAMMA, DELTA	4-node RAC members ALPHA, BETA, GAMMA, DELTA

Table 2. Oracle RAC host nodes

System Storage DS8000

Table 3 describes the two DS8000 images used in the backup and recovery exercises. This storage unit has two images and behaves as two separate DS8000 storage units The term logical unit number (LUN) and volume are used interchangeably on DS8000 to describe space allocations that are shared to host servers.

Role	**Primary storage (Site A)**	**Remote storage (Site B)**
Storage Unit	IBM.2107-75DF160	
Model	9B2	
Storage image Name	Image2	Image1
Storage ID	IBM.2107-75DF162	IBM.2107-75DF161
Code levels		
License Machine Code (LMC)	5.3.0.1136	5.3.0.1136
Storage Manager	5.1.600.196	5.1.600.196
DSCLI	6.1.3.20071026.1	6.1.3.20071026.1
WWNN	5005076307FFCF7D	5005076307FFC77D
RAID type	Raid 5 (7+P)	Raid 5 (7+P)
DDM capacity	300 GB	300 GB

Table 3. System Storage DS8000

DS8000 logical configuration

Volume configuration for Primary DS8000 image 2 (Site A) is described in Table 4. When using consistency groups for FlashCopy or Metro Mirror functions, the I/O is frozen on all volumes in the entire logical subsystem (LSS). Therefore, when creating the volumes, put them in separate LSSs based on ASM diskgroups. In these examples, the distribution of volumes to LSSs is as follows:

LSS 60 = Cluster Services volumes and backup volume
LSS 62 = DATA_DG_8K
LSS 64 = LOGS_DG_8K
LSS 66 = ARCH_DG_8K
LSS 35 FlashCopy targets

The FlashCopy target volumes are not visible to the primary RAC nodes.

ASM Diskgroup name	Primary Source		FlashCopy Target		LUN Size
	Source LUN ID	Source Name	FC_Target LUN ID	FC_Target Name	
DATA_DG_8K	6220	DATA_6220	3520	FC_DATA_3520	100 GB
	6221	DATA_6221	3521	FC_DATA_3521	100 GB
	6222	DATA_6222	3522	FC_DATA_3522	100 GB
LOGS_DG_8K	6430	LOGS_6430	3530	FC_LOGS_3530	10 GB
	6431	LOGS_6431	3531	FC_LOGS_3531	10 GB
	6432	LOGS_6432	3532	FC_LOGS_3532	10 GB
ARCH_DG_8K	6640	ARCH_6640	3540	FC_ARCH_3540	100 GB
	6641	ARCH_6641	3541	FC_ARCH_3541	100 GB
	6642	ARCH_6642	3542	FC_ARCH_3542	100 GB

Table 4. DS8000 primary RAC database volume configuration

Volume allocation for both Clone DS8000 storage volumes on image 2 and Remote DS8000 storage volumes on image 1 (Site B) are defined in Table 5. Assign the Clone volumes to the Clone RAC nodes.

	Volumes on SiteA		Volumes on Site B		
ASM Diskgroup	Source LUN ID	Clone ID	MM LUN ID	MM volume Name	Size
DATA_DG_8K	6220	7320	7320	MM_DATA_7320	100 GB
	6221	7321	7321	MM_DATA_7321	100 GB
	6222	7322	7322	MM_DATA_7322	100 GB
LOGS_DG_8K	6430	7530	7530	MM_LOGS_7530	10 GB
	6431	7531	7531	MM_LOGS_7531	10 GB
	6432	7532	7532	MM_LOGS_7532	10 GB
ARCH_DG_8K	6640	7740	7740	MM_ARCH_7740	100 GB
	6641	7741	7741	MM_ARCH_7741	100 GB
	6642	7742	7742	MM_ARCH_7742	100 GB

Table 5. DS8000 Clone and Remote RAC volume configuration

Oracle software installation

The Oracle binaries are installed creating three separate Oracle homes as shown in Table 6.

Clusterware home	CRS_HOME=/u01/crs/product/10.2/crs
ASM home	ORACLE_HOME=/u01/asm/product/10.2/asm
Database home	ORACLE_HOME=/u01/db/product/10.2/db

Table 6. Directories for Oracle binaries

Oracle Clusterware disks

When installing Oracle Clusterware, follow the Oracle installation documentation for Clusterware and define five 1 GB volumes: two for OCR and three for Vote. Ensure that each host node has the shared volumes identified consistently across all nodes in the cluster. These are labeled Shared-cluster services disks in Table 7.

Configuring primary Oracle RAC servers

Table 7 shows the storage allocations for each of the host nodes in the cluster.

Host-node name	**ALPHA**	**BETA**	**GAMMA**	**DELTA**
Shared-cluster services disks	2 OCR 3 Vote	2 OCR 3 Vote	2 OCR 3 Vote	2 OCR 3 Vote
Shared-data disks	9	9	9	9
Backup disk mounted on one node at a time	1	1	1	1
Multipath driver (same driver on all nodes)	Device mapper, or SDD PCM	Device mapper, or SDD PCM	Device mapper, or SDD PCM	Device mapper, or SDD PCM
Oracle Stack configured on local host disk		CRS_HOME /u01/crs/product/10.2/crs ASM_HOME /u01/asm/product/10.2/asm ORACLE_HOME /u01/db/product/10.2/db		
+ASM				
Instance name per host	+ASM1	+ASM2	+ASM3	+ASM4
Database name PRIME8K				
Database instance name per host	PRIME8K1	PRIME8K2	PRIME8K3	PRIME8K4

Table 7. Server nodes for 4-node cluster

Volume allocation for Oracle ASM and Oracle Database instances

On the Primary RAC nodes, create the ASM diskgroups: DATA_DG_8K, LOGS_DG_8K, ARCH_DG_8K. Assign disks to the ASM diskgroups and populate with database components, as shown in Table 8.

LUN name	**ASM diskgroup**	**ASM disk name (alias on host)**	**Contains database components**	**Additional directories**
DATA_6220	DATA_DG_8K	DS8K_DATA_1	Datafiles	
DATA_6221		DS8K_DATA_2		
DATA_6222		DS8K_DATA_3		
LOGS_6430	LOGS_DG_8K	DS8K_LOGS_1	redo logs controlfile (#2,#3)	
LOGS_6431		DS8K_LOGS_2		
LOGS_6432		DS8K_LOGS_3		
ARCH_6640	ARCH_DG_8K	DS8K_ARCH_1	Archived redo logs Spfile controlfile (#1)	+ARCH_DG/PRIME8K /rman_backups
ARCH_6641		DS8K_ARCH_2		
ARCH_6642		DS8K_ARCH_3		

Table 8. ASM diskgroup configuration and contents

RMAN and FlashCopy

The procedures describe how to use FlashCopy and Metro Mirror for backing up the database. RMAN is only used to make backup copies of the archive logs, controlfiles and spfiles to either a disk that is mountable to any node of the Oracle RAC cluster or to a directory in an Oracle ASM diskgroup. The datafile backup is handled with Advanced Copy Services.

The cloning scenario provides the ability to also backup the database to other media through RMAN. This could be done by bringing the database up to the mount state on the clone or remote site by using a backup controlfile, connecting to RMAN, and backing up the database (for example: database datafiles, controlfiles, spfiles, and archived logs). This process is not covered.

RMAN Recovery Catalog

Create an RMAN recovery catalog on a separate server. Then, register the database in the catalog. The recovery catalog database used in these examples is called *utils***.** User *rman* (with the password of *opensezme*) owns the recovery catalog. Use the following RMAN commands to connect to the catalog and register the database.

Command example 1: Connecting to RMAN catalog

```
[oracle@ALPHA ~]$ rman target / catalog rman/opensezme@utils;
RMAN> register database;
```

After registering the database, run the *report schema* command to view the registration information in the RMAN catalog output. Note the DBID for the database.

Command example 2: RMAN report schema

```
[oracle@ALPHA ~]$ rman target / catalog rman/opensezme@utils;

Recovery Manager: Release 10.2.0.4.0 - Production on Tue Aug 5 15:55:23 2008
Copyright (c) 1982, 2007, Oracle.  All rights reserved.
connected to target database: PRIME8K (DBID=3095618455)
connected to recovery catalog database
RMAN> register database;
database registered in recovery catalog
starting full resync of recovery catalog
full resync complete
RMAN> report schema;
Report of database schema

List of Permanent Datafiles
===========================
File Size(MB) Tablespace RB segs Datafile Name
---- -------- -------------------- ------- ------------------------
1    480      SYSTEM       YES  +DATA_DG_8K/prime8k/datafile/system.264.661895
2    40       UNDOTBS1     YES  +DATA_DG_8K/prime8k/datafile/undotbs1.265.661895
3    310      SYSAUX       NO  +DATA_DG_8K/prime8k/datafile/sysaux.266.661895
4    5        USERS        NO  +DATA_DG_8K/prime8k/datafile/users.263.661895
5    25       UNDOTBS2     YES +DATA_DG_8K/prime8k/datafile/undotbs2.262.661895
6    25       UNDOTBS3     YES +DATA_DG_8K/prime8k/datafile/undotbs3.260.661895
7    50       UNDOTBS4     YES +DATA_DG_8K/prime8k/datafile/undotbs4.267.661895
List of Temporary Files
=======================
File Size(MB) Tablespace  Maxsize(MB) Tempfile Name
---- -------- -------------------- ----------- --------------------
1    20       TEMP        32767  +DATA_DG_8K/prime8k/tempfile/temp.261.661895
```

RMAN backup directories

For this set of scenarios, a disk that contains the backup directory is mounted on the Oracle RAC node performing the backup. When recovering a database from this backup disk, it is important to mount the backup disk to the RAC node that performs the recovery. These examples use a disk that is mounted on the server that performs the backup with the directory structure:

/u01/backup/PRIME8K/ - spfile, controlfile and archived logs

This example also uses a directory in an ASM diskgroup:

+ARCH_DG_8K/PRIME8K/rman_backups - spfile, controlfile

Command example 3: Create backup directory in ASM diskgroup

```
[oracle@ALPHA ~]$ export ORACLE_HOME=/u01/asm/product/10.2/asm
[oracle@ALPHA ~]$ asmcmd -p
ASMCMD [+] > mkdir +ARCH_DG_8K/PRIME8K/rman_backups
```

RMAN backup controlfile

Controlfile backup types are illustrated in the following Command examples.

Command example 4: Backup controlfile to trace

```
sql 'alter database backup controlfile to trace';
```

Command example 5: Backup controlfile as a copy to the backup directory

```
backup as copy
 format '/u01/backup/PRIME8K/ctl_copy.ctl'
 current controlfile;
```

Command example 6: Backup controlfile as a RMAN dataset

```
backup
 format '/u01/backup/PRIME8K/cf_t%t_s%s_p%p'
 current controlfile;
```

Command example 7: Backup controlfile as a RMAN dataset to ASM

```
backup
 format '+ARCH_DG/PRIME8K/rman_backups/cf_t%t_s%s_p%p'
 current controlfile;
```

RMAN backup spfile

The spfile can be dumped to a regular text file or to a RMAN dataset with the following Command examples.

Command example 8: Create text init.ora file from spfile.

```
sql "create pfile=''/u01/backup/PRIME8K/init.ora.back'' from spfile";
```

Command example 9: Backup spfile as RMAN dataset to backup directory.

```
backup
 format '/u01/backup/PRIME8K/spfile_t%t_s%s_p%p'
 spfile;
```

Command example 10: Backup spfile as a RMAN dataset to ASM

```
backup
  format '+ARCH_DG_8K/PRIME8K/rman_backups/spfile_t%t_s%s_p%p'
  spfile;
```

RMAN archive log backup

This is shown as a as standard practice. No archive logs are restored for the purposes of the scenarios in this paper.

Command example 11: Backup archived logs to the backup directory

```
backup
  filesperset 10
  format '/u01/backup/PRIME8K/al_t%t_s%s_p%p'
  (archivelog all);
```

Configuration of Clone or Remote Oracle RAC nodes

Preparation of Clone or Remote Oracle RAC requires you to perform the following series of tasks:

1. Oracle RAC nodes should have the same operating-system level.
2. The DS8000 volumes must be allocated for Oracle Clusterware before installation of binaries.
3. Installation of Oracle binaries should be at the same patch level and in the same directory paths as the primary Oracle RAC for the purposes of production cloning.
4. Cluster services is configured.
5. Volumes allocated from DS8000 for data in the same number and size as the primary cluster's volumes

Table 9 shows the component requirements for the Clone and Remote cluster configuration.

Host-node name	**CL_ALPHA**	**CL_BETA**	**CL_GAMMA**	**CL_DELTA**
Shared-cluster services disks	2 OCR 3 Vote	2 OCR 3 Vote	2 OCR 3 Vote	2 OCR 3 Vote
Shared-data disks that are equal in number and size to a primary cluster	9	9	9	9
Backup disk mounted on one node at a time that is equal in size to the primary cluster backup disk	1	1	1	1
Multipath driver (same driver on all nodes)	Device mapper, or SDD PCM	Device mapper, or SDD PCM	Device mapper, or SDD PCM	Device mapper, or SDD PCM
Oracle stack configured on local host disk		CRS_HOME /u01/crs/product/10.2/crs ASM_HOME /u01/asm/product/10.2/asm ORACLE_HOME /u01/db/product/10.2/db		
+ASM				
Instance name per host	+ASM1	+ASM2	+ASM3	+ASM4
Database name				
Database instance name per host				

Table 9. Clone and Remote Oracle RAC node configuration

Oracle ASM diskgroup configuration

Table 10 contains the disk configuration for the clone and remote disks The Site B DS8000 volumes are allocated following the same pattern of separate LSSs for each ASM diskgroup.

Site A DS8000	**ASM diskgroup**	**ASM disk name (alias on host)**	**ASM diskgroup**	**Site B DS8000**
DATA_6220	DATA_DG_8K	DS8K_DATA_1	DATA_DG_8K	DATA_7320
DATA_6221		DS8K_DATA_2		DATA_7321
DATA_6222		DS8K_DATA_3		DATA_7322
LOGS_6430	LOGS_DG_8K	DS8K_LOGS_1	LOGS_DG_8K	LOGS_7530
LOGS_6431		DS8K_LOGS_2		LOGS_7531
LOGS_6432		DS8K_LOGS_3		LOGS_7532
ARCH_6640	ARCH_DG_8K	DS8K_ARCH_1	ARCH_DG_8K	ARCH_7740
ARCH_6641		DS8K_ARCH_2		ARCH_7741
ARCH_6642		DS8K_ARCH_3		ARCH_7742

Table 10. ASM diskgroup configuration for Clone Oracle RAC cluster

Preparation of Clone or Remote Oracle RAC

Set up the cluster by duplicating the Oracle binaries at the same version and patch level as is the case at the primary site. The directory paths for the Oracle binary installations should be the same as at the primary site. DS8000 LUNs need to be allocated and configured for the Shared Oracle Clusterware disks, as specified in Table 9. Data disks must be allocated such that they are equal in number and size to the primary Oracle RAC.

Clone or Remote Oracle ASM setup

Preparation of Clone and Remote Oracle ASM diskgroup setup is as follows.

1. Use Oracle Database Configuration Assistant (DBCA) to set up the ASM instances on the Clone or Remote Oracle RAC. Do not define ASM diskgroups.
2. Use the Oracle srvctl tool to stop ASM instances on all target Oracle RAC nodes.
3. On node 1 of the Clone or Remote cluster, set the asm_diskgroups parameter such that it is equal to the primary cluster pfile asm_diskgroups parameter. For example, on this setup, the parameter is set as:
   asm_diskgroups='DATA_DG_8K','LOGS_DG_8K','ARCH_DG_8K'.
4. Copy this init.ora to the remaining Clone or Remote nodes.

Clone or Remote database instance setup

There are four important files to copy or edit on the Clone or Remote Oracle RAC nodes to migrate the database from the primary: initPRIMEx.ora, orapwPRIMEx, oratab file, and tnsnames.ora:

1. Copy the initPRIMEx.ora file (found at /u01/db/product/10.2/db/dbs/initPRIMEx.ora), which contains a pointer to the spfile in the Oracle ASM.
2. Copy the orapwPRIMEx file (found at /u01/db/product/10.2/dbs/orapwPRIMEx), which is the database password file.
3. Check the oratab file and add the following (the syntax of +ASMx is dependent on the node):

   ```
   +ASMx:/u01/asm/product/10.2/asm:N
   PRIME8K:/u01/db/product/10.2/db:N
   ```

4. Copy the tnsnames.ora file from the primary nodes' subdirectory /u01/db/product/10.2/db/network/admin/. Then, edit it by replacing the original host-node names with the Clone or Remote Oracle RAC host names.
5. Create subdirectories for database logs and trace files under /u01/admin/PRIME (that is, adump, bdump, cdump, dpdump, hdump, pfile and udump).
6. After editing the files for the first node, copy initPRIMEx.ora, orapwPRIMEx, tnsnames.ora and the oratab files to the remaining nodes in the Clone or Remote Oracle RAC node, while being sure to also change the x value to the appropriate node number.

DS8000 Advanced Copy Services overview and configuration

The DS8000 Advanced Copy Services functions FlashCopy and Metro Mirror processes are described in full in the following section.

FlashCopy

Base FlashCopy sequence of events

When you set up a FlashCopy operation:

1. A relationship is established between the source and target volumes.
2. A bitmap of the source volume is created when the relationship is created.
3. Once relationship and bitmap are created, the target volume can be accessed as though all the data had been physically copied.
4. While a relationship between the source and target volume exists, a background process copies the tracks from the source to the target volume.

Refresh target volume resyncflash

Resyncflash provides the ability to refresh a target volume in a FlashCopy relationship as long as the *-record* and *-persist* parameters were used when creating the relationship. When a subsequent FlashCopy operation is initiated through *resyncflash*, only the tracks changed on both the source and target need to be copied from the source to the target. The direction of the refresh can also be reversed.

resyncFlashCopy operations

1. Full FlashCopy is created with the -record and -persistent option.
2. The change recording bitmaps are used for recording the tracks that are changed on the source and target volumes after the last FlashCopy.
3. After creating the change recording bitmaps, Copy Services records the information for the updated tracks to the bitmaps. The FlashCopy relationship persists even if all of the tracks have been copied from the source to the target.
4. With the next resyncflash, Copy Services checks the change recording bitmaps and copies only the changed tracks to the target volumes. If some tracks on the target volumes are updated, these tracks are overwritten by the corresponding tracks from the source volume.

Reverse FlashCopy - reverseflash

Reverse FlashCopy is used to restore, using the target disk as the new source. The reverse restore operation cannot be done unless the background copy in the original direction is finished.

1. Full FlashCopy is created with the -record -persistent option.
2. The change recording bitmaps are used for recording the tracks that are changed on the source and target volumes after the last FlashCopy.
3. After creating the change recording bitmaps, Copy Services records the information for the updated tracks to the bitmaps. The FlashCopy relationship persists, even if all of the tracks have been copied from the source to the target.
4. Issue *reverseflash* command with parameters, Copy Services checks the change recording bitmaps and copies modified tracks from the target volume to the source volume. The background copy process must be complete before you can reverse the order of the FlashCopy relationship to its original source and target relationship.

Consistency for a group of mappings

ASM diskgroups have data spread over multiple volumes. Therefore, all member volumes in the ASM diskgroup must be copied based on the same point in time. FlashCopy requires the -freeze option to create a consistent PiT copy across multiple volumes within the group. In the examples shown, a sequence number is assigned to all members of the consistency group. When the relationship is reversed, the sequence number is changed to show that the members are in a reversed relationship.

With the Freeze FlashCopy Consistency Group option, the DS8000 holds off I/O activity to a volume for a time, by putting the source volume in a queue full state. The cessation of I/O activity to the volumes in a consistency group is visible as an I/O freeze at the Oracle level.

Therefore, a time slot can be created during which dependent write updates do not occur, and FlashCopy uses that time slot to obtain a consistent PiT copy of the related volumes. I/O activity resumes when all FlashCopy replicas are established. When using the freeze parameter to create consistency groups, the source volumes are put into a queue full state freezing I/O, reset the queue full state by issuing the unfreezeflash command on the LSS immediately following FlashCopy command.

Example FlashCopy commands and parameters

The commands used for the scenarios were taken directly from Dscli help. Table 11 includes the FlashCopy common parameters a description for each parameter is provided. Any parameters that are not found in this table are defined with the appropriate command.

parameter	Action indicated by parameter
- dev	storage_image_ID This parameter is required if you do not specify a fully qualified ID for the source in the dscli profile.
-record	Records the changed tracks on both volumes within a FlashCopy pair. Select this parameter when you establish an initial FlashCopy volume pair that you intend to use with the resyncflash command. The -persist parameter is automatically designated when the -record parameter is selected.
-persist	Relationship exists until explicitly removed
-freeze	Establish a Consistency Group ensuring member volumes are consistent at the same PIT. Triggers the queue full condition for the source volumes. All writes to the source volumes are queued by the host and are written after the queue full condition is reset (unfreezeflash LSS.) During the queue full condition, the source volume reports long busy status. The queue full condition is reset by an extended long busy timeout. The timeout condition affects all FlashCopy source volumes that are contained within a respective logical subsystem and that are established or modified with the -freeze parameter.
-tgtpprc	Allows the FlashCopy target volume to be a Remote Mirror and Copy source volume.
-seqnum	Associates the FlashCopy relationships that are established with the specified sequence number. This sequence number can be used as an identifier for a relationship or group of relationships. Only the relationships that are modified successfully by the resyncflash command get the specified sequence number, leaving the ones that failed with the previous one (if they were previously specified).
source:target:	Increments a FlashCopy relationship for the source and target volume pairs with the IDs specified. A FlashCopy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. Separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume. This can be a list of pairs as in 6020:7520 6021:7521 or as 6020-6022:7520-7522
-l	For ls commands: Displays the default output plus copy indicator, out of sync tracks, date created, and date synchronized.
-fmt delim -delim |	Formats the output using the bar as a delimiter.

Table 11. FlashCopy common parameters

chlss

The chlss command modifies a logical subsystem (LSS). The feature of creating consistency group freezes I/O on the source logical subsystem.

Parameters

-pprcconsistgrp enable | disable

Enables a volume that is associated with a logical subsystem to become suspended and enter an extended long busy state if it has not received a notification that a consistency group has been created. Otherwise, the volumes associated with the LSS do not go to a long-busy state.

Command example 12: Chlss to enable consistency groups timeout

```
chlss -pprcconsistgrp enable  62
chlss -pprcconsistgrp enable  64
chlss -pprcconsistgrp enable  66
```

Create FlashCopy relationships

Create the pairings of source and target volumes on the DS8000. In this example, mappings are placed in a consistency grouping with the -freeze parameter. In the pairings, source volumes are designated as volume A and target volumes are designated as volume B for each member pair. For example, source volume A is LUN ID 6220 and target volume B is LUN ID 3520.

Command example 13: Mkflash to create volume pairings

DATA_DG_8K volumes 6220, 6221, 6222

```
mkflash -persist -record -freeze -seqnum 20 6220:3520 6221:3521 6222:3522
```

LOGS_DG_8K volumes 6430, 6431, 6432

```
mkflash -persist -record -freeze -seqnum 30 6430:3530 6431:3531 6432:3532
```

ARCH_DG_8K volumes 6640, 6641, 6642

```
mkflash -persist -record -freeze -seqnum 40 6640:3540 6641:3541 6642:3542
```

Unfreezing the LSS

When the -freeze parameter is used to create consistency groups, a queue full condition is placed on all volumes in the LSS, whether they are members of the current FlashCopy pairings or not. The unfreezeflash command resets the queue full condition established with the -freeze parameter when the mkflash or resyncflash commands are issued.

Command example 14: Unfreeze the logical subsystem to reset the long busy timeout

```
dscli> unfreezeflash 62 64 66
```

View the FlashCopy status

The lsflash command displays a list of FlashCopy relationships and status information for each FlashCopy relationship in the list.

Command example 15: List FlashCopy relationships and status

```
lsflash -l -fmt delim -delim | 6200-6642
```

Table 12 lists and describes the fields that are important to monitor the FlashCopy process.

ID	Specifies the FlashCopy pair ID. This ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. Ex.6020:7520
Sequence Num	Sequence number that is associated with the FlashCopy relationship.
ActiveCopy	(Enabled or Disabled) whether the background copy process is currently active for this FlashCopy relationship.
Recording	(Enabled or Disabled) whether this FlashCopy relationship was established with the record changes option.
Persistent	(Enabled or Disabled) whether this FlashCopy relationship was established with the persistent option.
OutOfSyncTracks	The number of tracks that are not synchronized for this FlashCopy relationship. The maximum value that can be displayed is dependent on the source volume size.
DateCreated	Date and the time that the FlashCopy relationship was established.
DateSynced	Date and time this FlashCopy relationship was synchronized, or null (-) if the relationship is not synchronized.

Table 12. Report fields descriptions for lsflash

Results of lsflash command

```
dscli> dscli> lsflash -fmt delim -delim | 6200-6642
Date/Time: August 5, 2008 8:32:02 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID|SrcLSS|SequenceNum|Timeout|ActiveCopy|Recording|Persistent|Revertible|Sourc
eWriteEnabled|TargetWriteEnabled|BackgroundCopy
==============================================================================
6220:3520|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6221:3521|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6222:3522|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6430:3530|64|30|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6431:3531|64|30|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6432:3532|64|30|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6640:3540|66|40|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6641:3541|66|40|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6642:3542|66|40|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
```

In the case of a single flash copy, you can use the stanza format for output, such as the following:

```
dscli> lsflash -l -fmt stanza 6222
Date/Time: August 5, 2008 8:32:51 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID                 6222:3522
SrcLSS             62
SequenceNum        20
Timeout            60
ActiveCopy         Disabled
Recording          Enabled
Persistent         Enabled
Revertible         Disabled
SourceWriteEnabled Enabled
TargetWriteEnabled Enabled
BackgroundCopy     Enabled
OutOfSyncTracks    0
DateCreated        Mon Aug 04 21:20:20 PDT 2008
DateSynced         Mon Aug 04 21:20:20 PDT 2008
State              Valid
AllowTgtSE         Disabled
```

Refreshing the FlashCopy with resyncflash

The resyncflash command increments an existing FlashCopy pair that has been established with the *-record* and *-persist* parameters. When a pair is established with the *-record* and *-persist* parameters, the pair initially synchronizes and then a record of all host write operations to the source is maintained in the source volumes. When the resyncflash command is issued on the pair, the changed data that is written to the source is copied to the target.

Command example 16: Increment the FlashCopy with resyncflash

```
resyncflash -persist -record -freeze -seqnum 20 6220-6222:3520-3522
resyncflash -persist -record -freeze -seqnum 30 6430-6432:3530-3532
resyncflash -persist -record -freeze -seqnum 40 6640-6642:3540-3542
```

Return output example:

```
dscli> resyncflash -persist -record -freeze -seqnum 20 6220-6222:3520-3522
Date/Time: August 5, 2008 8:34:36 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
CMUC00168I resyncflash: FlashCopy volume pair 6220:3520 successfully
resynchronized.
CMUC00168I resyncflash: FlashCopy volume pair 6221:3521 successfully
resynchronized.
CMUC00168I resyncflash: FlashCopy volume pair 6222:3522 successfully
resynchronized.CMUC00168I resyncflash: FlashCopy volume pair 6222:3522
successfully resynchronized.
```

When a -freeze parameter is used, you should immediately follow up with an *unfreezeflash* on the LSS.

```
dscli> unfreezeflash 62
Date/Time: August 5, 2008 8:34:55 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
CMUC00172I unfreezeflash: FlashCopy consistency group for logical subsystem
62: successfully reset.
```

Resulting output from *lsflash* command after resynchronizing:

```
dscli> lsflash -l -fmt delim -delim | 6220-6222
Date/Time: August 5, 2008 8:37:34 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID|SrcLSS|SequenceNum|Timeout|ActiveCopy|Recording|Persistent|Revertible|Sourc
eWriteEnabled|TargetWriteEnabled|BackgroundCopy|OutOfSyncTracks|DateCreated|Da
teSynced|State|AllowTgtSE
================================================================================
================================================================
==================================
6220:3520|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled|0
|Mon Aug 04 21:20:20 PDT 2008|Tue Aug 05 08:32:52 PDT 2008|Valid|Disabled
6221:3521|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled|0
|Mon Aug 04 21:20:20 PDT 2008|Tue Aug 05 08:32:52 PDT 2008|Valid|Disabled
6222:3522|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled|0
|Mon Aug 04 21:20:20 PDT 2008|Tue Aug 05 08:32:52 PDT 2008|Valid|Disabled
```

Changing the direction of the copy function with reverseflash

The direction of a FlashCopy relationship can be reversed, where the volume that was previously defined as the target (volume B) becomes the source for the volume that was previously defined as the source (volume A and is now the target). The data that has changed is copied to the volume that was previously defined as the source (volume A).

For example, you create a FlashCopy relationship between source volume A and target volume B. Data loss occurs on source volume A. To recover, reverse the FlashCopy relationship so that volume B is copied to volume A.

Change the value of the seqnum parameter to identify the relationship as being reversed. There is no *-freeze* parameter on *reverseflash*.

When the FlashCopy pairs were made, the relationship is 6220:3520 or A to B. When the *reverseflash* command is issued the relationship must be stated as reversing the A to B relationship.

Command example 17: Change direction of FlashCopy function

```
reverseflash -persist -record -tgtpprc -seqnum 22 6220-6222:3520-3522
reverseflash -persist -record -tgtpprc -seqnum 33 6430-6432:3530-3532
reverseflash -persist -record -tgtpprc -seqnum 44 6640-6642:3540-3542
```

Returns for each consistency group:

```
dscli> reverseflash -persist -record -tgtpprc -seqnum 22 6220-6222:3520-3522
Date/Time: August 27, 2008 9:09:06 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
CMUC00169I reverseflash: FlashCopy volume pair 6220:3520 successfully
reversed.
CMUC00169I reverseflash: FlashCopy volume pair 6221:3521 successfully
reversed.
CMUC00169I reverseflash: FlashCopy volume pair 6222:3522 successfully
reversed.
```

lsflash to show reversed FlashCopy pairs

Command example 18: List reversed pairs

```
lsflash -l -fmt delim -delim | -seqnum 22 3520-3522
```

Note: OutOfSyncTracks and source and target designations

```
dscli> lsflash -l -fmt delim -delim | -seqnum 22 3520-3522
Date/Time: August 27, 2008 9:10:57 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID|SrcLSS|SequenceNum|Timeout|ActiveCopy|Recording|Persistent|Revertible|Sourc
eWriteEnabled|TargetWriteEnabled|BackgroundCopy|OutOfSyncTracks|DateCreated|Da
teSynced|State|AllowTgtSE
================================================================================
==========================================================
3520:6220|35|22|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled|0
|Tue Aug 26 21:24:55 PDT 2008|Wed Aug 27 09:08:54 PDT 2008|Valid|Disabled
3521:6221|35|22|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled|0
|Tue Aug 26 21:24:55 PDT 2008|Wed Aug 27 09:08:54 PDT 2008|Valid|Disabled
3522:6222|35|22|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled|0
|Tue Aug 26 21:24:55 PDT 2008|Wed Aug 27 09:08:54 PDT 2008|Valid|Disabled
```

Restoring FlashCopy relationships to original direction with reverseflash

To restore the direction of the FlashCopy that has been reversed, an additional *reverseflash* is issued (reverse the reverse). This reestablishes the relationship to the original A to B relationship. Resetting the -seqnum to the original forward designation affirms the command was successful when you list the relationships. This process prepares the FlashCopy to be used in regular backup configuration. Ensure there is no Active Copy on any of the relationships before issuing this command.

Command example 19: Restore FlashCopy relationships to their original copy direction

```
reverseflash -persist -record -tgtpprc -seqnum 20 3520-3522:6220-6222
reverseflash -persist -record -tgtpprc -seqnum 30 3530-3532:6430-6432
reverseflash -persist -record -tgtpprc -seqnum 40 3540-3542:6640-6642
```

Note: Metro Mirror relationships were created on the source volumes, followed by the creation of FlashCopy relationships. *Reverseflash* requires the additional parameter *-tgtpprc* when both FlashCopy and Metro Mirror relationships are using the same source volumes.

```
dscli> lsflash -fmt delim -delim | 6220-6222
Date/Time: August 5, 2008 9:07:02 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID|SrcLSS|SequenceNum|Timeout|ActiveCopy|Recording|Persistent|Revertible|Sourc
eWriteEnabled|TargetWriteEnabled|BackgroundCopy
==============================================================================
==============================================
6220:3520|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6221:3521|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6222:3522|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
```

lsflash

Command example 20: Verify FlashCopy relationships

```
dscli> lsflash -fmt delim -delim | 6220-6642
Date/Time: August 5, 2008 9:07:02 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID|SrcLSS|SequenceNum|Timeout|ActiveCopy|Recording|Persistent|Revertible|Sourc
eWriteEnabled|TargetWriteEnabled|BackgroundCopy
==============================================================================
6220:3520|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6221:3521|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6222:3522|62|20|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6430:3530|64|30|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6431:3531|64|30|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6432:3532|64|30|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6640:3540|66|40|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6641:3541|66|40|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
6642:3542|66|40|60|Disabled|Enabled|Enabled|Disabled|Enabled|Enabled|Enabled
```

Metro Mirror

The following tasks show the sequence of events in configuring Metro Mirror relationships

1. Create storage unit relationships
2. List the available Fibre Channel links
3. Define the paths for each Metro Mirror consistency group
4. List the pprcpaths for the source LSS IDs
5. Make the remote mirror and copy relationship with *mkpprc* command
6. List PPRC relationships, noting status of Out of Sync Tracks and State
7. When the relationships state is full-duplex mode, you can now proceed

Example Metro Mirror commands and parameters

This section describes the configuration and implementation of Metro Mirror for this project.

Creating storage unit relationships

Metro Mirror Remote copy functions were formerly known as Peer to Peer Remote Copy (PPRC) functions. The PPRC commands encompass Metro Mirror, Global Mirror, and Global Copy. Metro Mirror is covered in this testing. I/O ports between storage units must be either direct connect between ports or facilitated through zoning on the switches used. Table 13 shows the detailed information required for the creation of the PPRC connections and relationships.

	PRIMARY Site A	**REMOTE Site B**
Storage	Jmt1 image2	Jmt1 image1
Serial	IBM.2107-75DF162	IBM.2107-75DF161
WWNN	5005076307FFCF7D	5005076307FFC77D
LSS pairings	**-srclss**	**-tgtlss**
-srclss on primary and -tgtlss on remote	62	73
	64	75
	66	77
I/O port	I0231	I0031
Volume consistency grouping	**Volume id**	**Maps to volume ID**
DATA	6220-6222	7320-7322
LOGS	6430-6432	7530-7532
ARCH	6640-6642	7740-7742

Table 13. Metro Mirror Site A and Site B relationship information

Table 14 shows the Metro Mirror common parameters, A description for each parameter (taken directly from the dscli help) is provided. Any parameters that are not found in this table are defined with the appropriate command.

-dev	Specifies the storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the source LSS IDs and source port IDs.
-remotedev	Specifies the ID of the secondary storage image, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified target port ID or if the -dev parameter is used.
-remotewwnn	Specifies the worldwide node name of the secondary storage image. The format is a 16-hexadecimal ID.
-srclss	Specifies the source logical subsystem ID.
-tgtlss	Specifies the target logical subsystem ID.
-consistgrp	Creates a consistency group for the remote mirror and copy volume pairs. A remote mirror and copy consistency group is a set of remote mirror and copy volume pairs that have the same source and target LSS.
source_port_ID: target_port_ID	The source and target ports must be fibre-channel I/O ports that are configured for point-to-point or switched fabric topology
-fullid	Give fully qualified names in the listing

Table 14. Metro Mirror common parameters

List the available Fibre Channel links

Verify the ports that are available for creating the paths.

The lsavailpprcport command displays a list of I/O ports that can be defined as remote mirror and copy paths.

Command example 21: List available pprc ports

```
lsavailpprcport -l -fullid -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -
remotewwnn 5005076307FFC77D 62:73

lsavailpprcport -l -fullid -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -
remotewwnn 5005076307FFC77D 64:75

lsavailpprcport -l -fullid -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -
remotewwnn 5005076307FFC77D 66:77
```

Example command result

```
dscli> lsavailpprcport -l -fullid -dev IBM.2107-75DF162 -remotedev IBM.2107-
75DF161 -remotewwnn 5005076307FFC77D 62:73
Date/Time: August 7, 2008 8:37:32 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
Local Port            Attached Port          Type Switch ID Switch Port
==========================================================================
IBM.2107-75DF162/I0231 IBM.2107-75DF161/I0031 FCP  NA        NA
IBM.2107-75DF162/I0301 IBM.2107-75DF161/I0101 FCP  NA        NA
```

Define the paths for each Metro Mirror consistency group

The mkpprcpath command establishes or replaces a remote mirror and copy path between source and target logical subsystems (LSSs) over a fibre-channel connection. Create one path for each consistency group from Site A to Site B.

Command example 22: Create pprc path from storage Site A to storage Site B

DATA_DG_8K

```
mkpprcpath -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -remotewwnn
5005076307FFC77D -consistgrp -srclss 62 -tgtlss 73 I0231:I0031
```

LOGS_DG_8K

```
mkpprcpath -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -remotewwnn
5005076307FFC77D -consistgrp -srclss 64 -tgtlss 75 I0231:I0031
```

ARCH_DG_8K

```
mkpprcpath -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -remotewwnn
5005076307FFC77D -consistgrp -srclss 66 -tgtlss 77 I0231:I0031
```

Command example 23: List PPRC path from storage Site A to storage Site B

```
dscli> lspprcpath 62 64 66
Date/Time: August 7, 2008 8:35:08 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
Src Tgt State   SS   Port  Attached Port Tgt WWNN
=======================================================
62  73  Success FF73 I0231 I0031         5005076307FFC77D
64  75  Success FF75 I0231 I0031         5005076307FFC77D
66  77  Success FF77 I0231 I0031         5005076307FFC77D
```

List the pprcpaths for the source LSS IDs

Displays the remote mirror and copy paths that are defined for the specified source LSS IDs:

Command example 24: List PPRC path for specific LSS IDs

```
lspprcpath -l -fmt delim -delim | -fullid 62 64 66 (from storage a)
lspprcpath -l -fmt delim -delim | -fullid 73 75 77 (from storage b)
```

```
dscli> lspprcpath -l -fmt delim -delim | -fullid 62 64 66
Date/Time: August 6, 2008 7:17:12 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
Src|Tgt|State|SS|Port|Attached Port|Tgt WWNN
=========================================
IBM.2107-75DF162/62|IBM.2107-75DF161/73|Success|FF73|IBM.2107-
75DF162/I0230|IBM.2107-75DF161/I0103|5005076307FFC2C8
IBM.2107-75DF162/64|IBM.2107-75DF161/75|Success|FF75|IBM.2107-
75DF162/I0230|IBM.2107-75DF161/I0103|5005076307FFC2C8
IBM.2107-75DF162/66|IBM.2107-75DF161/77|Success|FF77|IBM.2107-
75DF162/I0230|IBM.2107-75DF161/I0103|5005076307FFC2C8
```

Make the remote mirror and copy relationship with mkpprc

The *mkpprc* command establishes a remote mirror and copy (formerly PPRC) relationship for volume pairs.

Command example 25: Make Metro Mirror (PPRC)volume relationships

DATA_DG_8K 6220,6221,6222 map to 7320,7321,7322

```
mkpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type mmir -mode full
6220:7320 6221:7321 6222:7322
```

LOGS_DG_8K 6430,6431,6432 map to 7530,7531,7532

```
mkpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type mmir -mode full
6430:7530 6431:7531 6432:7532
```

ARCH_DG_8K 6640,6641,6642 map to 7740,7741,7742

```
mkpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type mmir -mode full
6640:7740 6641:7741 6642:7742
```

Additional Parameters

-type mmir: Metro Mirror maintains the remote mirror and copy relationship in a consistent (synchronous) manner by returning I/O write completion status to the application when the updates are committed to the target. This process becomes slower as the physical distance between source and target increases.

-mode full: Specifies the initial data copy mode for the remote mirror and copy volume relationships. Full mode copies the entire source volume to the target volume.

List PPRC relationships

The lspprc command displays a list of remote mirror and copy (formerly PPRC) volume relationships for a storage image, and status information for each remote mirror and copy volume relationship in the list.

Command example 26: List PPRC relationships

```
lspprc -l -fmt delim -delim | 6220:7320 6221:7321 6222:7322  6430:7530
6431:7531 6432:7532 6640:7740 6641:7741 6642:7742
or
lspprc -l -fmt delim -delim | 6220-6642
```

Resulting report field descriptions:

ID	Specifies the source and target volume IDs of a PPRC volume relationship.
State	Displays the current PPRC volume relationship state. One of the following values is displayed
Copy-Pending	Indicates that the relationship is copy pending. A Global Copy (Extended Distance) relationship is always copy pending
Full-Duplex	Indicates that the relationship is full duplex.
Suspended	Indicates that the relationship is suspended. The Reason attribute might indicate why the relationship is suspended
Target Copy Pending	Indicates that the source volume is unknown or cannot be queried and the target state is copy pending
Target Full-Duplex	Indicates that the source volume is unknown or cannot be queried and the target state is full duplex
Target Suspended	Indicates that the source volume is unknown or cannot be queried and the target state is suspended
Not Remote Copy Pair	Indicates that the relationship is not a remote copy pair.
Invalid-State	Indicates that the relationship state is invalid
Reason	Indicates why the remote copy and mirror volume relationship is suspended.
Type	Indicates that the remote copy and mirror volume relationship is a Metro Mirror (synchronous) relationship, a Global Copy (extended distance) relationship, or the relationship type is unknown.
Out Of Sync Tracks	Indicates the number of tracks that are not synchronized for this relationship. The maximum value is dependent on the source volume size.
Tgt Read	Indicates that Read IO operations to the target volume are allowed.
Src Cascade	Indicates that the source volume of this relationship is enabled to also be a target volume of a different relationship.
Tgt Cascade	Indicates that the target volume of this relationship is enabled so that it is also a source volume for a different relationship.
DateSuspended	Indicates the date when this relationship was last suspended.
SourcerLSS	Indicates the Consistency Group LSS ID that is associated with the source volume of this PPRC volume relationship.
Timeout (secs)	Indicates the Consistency Group Long Busy Timeout setting for the LSS ID that is associated with the source volume of this PPRC volume relationship. The timeout value is either the default value of 120 seconds, or a user-specified value of 1 to 600 000 seconds.
Crit mode	Indicates whether the remote copy and mirror primary volume represents a critical volume.
First pass status	Indicates the first pass Global Copy state of the source volume. The data is valid only when you query the primary box and the queried pair is a Global Copy pair.

Table 15. Report field descriptions for lspprc

List PPRC relationships, noting status of Out of Sync Tracks and State

Command example 27: List PPRC following the creation of the relationships

```
dscli> lspprc -l -fmt delim -delim | 6220-6222
Date/Time: August 7, 2008 7:32:03 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID|State|Reason|Type|Out Of Sync Tracks|Tgt Read|Src Cascade|Tgt Cascade|Date
Suspended|SourceLSS|Timeout (secs)|Critical Mode|First Pass Status
==========================================================================
==================================================================
6220:7320|Copy Pending|-|Metro Mirror|1442644|Disabled|Disabled|invalid|-
|62|300|Disabled|Invalid
6221:7321|Copy Pending|-|Metro Mirror|1441884|Disabled|Disabled|invalid|-
|62|300|Disabled|Invalid
6222:7322|Copy Pending|-|Metro Mirror|1561045|Disabled|Disabled|invalid|-
|62|300|Disabled|Invalid
```

Lspprc: Note the State of the relationships is now Full Duplex, the Out Of Sync Tracks is zero; the volumes are updated synchronously.

```
dscli> lspprc -l -fmt delim -delim | 6220-6222
Date/Time: August 27, 2008 9:20:42 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID|State|Reason|Type|Out Of Sync Tracks|Tgt Read|Src Cascade|Tgt Cascade|Date
Suspended|SourceLSS|Timeout (secs)|Critical Mode|First Pass Status|Incremental
Resync|Tgt Write|GMIR CG|PPRC CG|AllowTgtSE|DisableAutoResync
==========================================================================
==========================================================================
==================================================================
6220:7320|Full Duplex|-|Metro Mirror|0|Disabled|Disabled|Invalid|-
|62|60|Disabled|Invalid|Disabled|Disabled|Disabled|Enabled|Disabled|-
6221:7321|Full Duplex|-|Metro Mirror|0|Disabled|Disabled|Invalid|-
|62|60|Disabled|Invalid|Disabled|Disabled|Disabled|Enabled|Disabled|-
6222:7322|Full Duplex|-|Metro Mirror|0|Disabled|Disabled|Invalid|-
|62|60|Disabled|Invalid|Disabled|Disabled|Disabled|Enabled|Disabled|-
```

Listing the PPRC pairs from Site B

```
dscli> lspprc -l -fmt delim -delim | 7320-7742
Date/Time: August 7, 2008 8:32:03 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF161
ID|State|Reason|Type|Out Of Sync Tracks|Tgt Read|Src Cascade|Tgt Cascade|Date
Suspended|SourceLSS|Timeout (secs)|Critical Mode|First Pass Status

==================================================================
6220:7320|Target Copy Pending|-|Metro Mirror|0|Disabled|invalid|Disabled|-
62|unknown|Disabled|Invalid
6221:7321|Target Copy Pending|-|Metro Mirror|0|Disabled|invalid|Disabled|-
62|unknown|Disabled|Invalid
6222:7322|Target Copy Pending|-|Metro Mirror|0|Disabled|invalid|Disabled|-
62|unknown|Disabled|Invalid
```

Using Metro Mirror consistency groups

The following steps show an overview of the process to create a consistency group copy on a remote storage device.

1. Freezepprc: creating consistent copies across group members, immediately followed by an unfreezepprc.
2. Failoverpprc: enables read-write at the remote site.
3. Failbackpprc: Site A to Site B resumes the Metro Mirror copy from the primary site.

Metro Mirror failover failback process

These steps are required to perform the failover and failback process in which a failback is done from the local site Site A to the remote site Site B. This allows you to bring up a test database on the remote volumes, or backup the remote database to tape. Then, when the remote processing is complete, Copy Services resynchronizes the remote volumes from the local (production) volumes by copying only changed tracks from the local to the remote volumes.

Freezing the relationships pprcfreeze

The *pprcfreeze* command queues updates to the Site A volumes in Metro Mirror relationships across the affected LSSs. This process ensures that the Site B volumes are consistent at the time of the freeze. (One command per LSS pair is required.) As a result of the freeze action, the following processing occurs:

1. I/O processing to the Metro Mirror volume pairs is temporarily queued during the time that updates are frozen.
2. The volume pairs that are associated with the source and target LSSs are suspended. During this time, updates are collected using the change recording feature on the Site A volumes.
3. The established paths between the LSS pairs are disabled (removed).

Command example 28: Make the consistent copy groupings freezepprc

DATA_DG_8K

```
freezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 62:73
```

LOGS_DG_8K

```
freezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 64:75
```

ARCH_DG_8K

```
freezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 66:77
```

Resume I/O activities on Site A following a freeze

Issue the *unfreezepprc* command to allow I/O activity to resume for the specified volume pairs.

Command example 29: Allow I/O to resume on source with unfreezepprc

DATA_DG_8K

```
unfreezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 62:73
```

LOGS_DG_8K

```
unfreezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 64:75
```

ARCH_DG_8K

```
unfreezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 66:77
```

Display PPRC status on Site A and Site B

List the status of the PPRC pairs after *freezepprc* on Site A. Note the State is Suspended and Reason is Freeze.

```
dscli> lspprc -l -fmt delim -delim | 6220-6222
Date/Time: August 27, 2008 9:26:05 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF162
ID|State|Reason|Type|Out Of Sync Tracks|Tgt Read|Src Cascade|Tgt Cascade|Date
Suspended|SourceLSS|Timeout (secs)|Critical Mode|First Pass Status|
Incremental Resync|Tgt Write|GMIR CG|PPRC CG|AllowTgtSE|DisableAutoResync
==========================================================================
==========================================================================
============================================================
6220:7320|Suspended|Freeze|Metro Mirror|0|Disabled|Disabled|Invalid|-
|62|60|Disabled|Invalid|Disabled|Disabled|Disabled|Enabled|Disabled|-
6221:7321|Suspended|Freeze|Metro Mirror|0|Disabled|Disabled|Invalid|-
|62|60|Disabled|Invalid|Disabled|Disabled|Disabled|Enabled|Disabled|-
6222:7322|Suspended|Freeze|Metro Mirror|0|Disabled|Disabled|Invalid|-
|62|60|Disabled|Invalid|Disabled|Disabled|Disabled|Enabled|Disabled|-
```

List the status of the PPRC pairs after *freezepprc* on Site B. Note the State is Target Full Duplex and Reason is null.

```
dscli> lspprc -l -fmt delim -delim | 7320-7322
Date/Time: August 27, 2008 9:34:25 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF161
ID|State|Reason|Type|Out Of Sync Tracks|Tgt Read|Src Cascade|Tgt Cascade|Date
Suspended|SourceLSS|Timeout (secs)|Critical Mode|First Pass Status|
Incremental Resync|Tgt Write|GMIR CG|PPRC CG|AllowTgtSE|DisableAutoResync
==========================================================================
==========================================================================
============================================================
6220:7320|Target Full Duplex|-|Metro Mirror|0|Disabled|Invalid|Disabled|-
|62|unknown|Disabled|Invalid|Disabled|Disabled|Disabled|Disabled|Disabled|-
6221:7321|Target Full Duplex|-|Metro Mirror|0|Disabled|Invalid|Disabled|-
|62|unknown|Disabled|Invalid|Disabled|Disabled|Disabled|Disabled|Disabled|-
6222:7322|Target Full Duplex|-|Metro Mirror|0|Disabled|Invalid|Disabled|-
|62|unknown|Disabled|Invalid|Disabled|Disabled|Disabled|Disabled|Disabled|-
```

At Site B (remote site), issue a failover command to the B to A volume pairs

The *failoverpprc* command changes the secondary device into a primary suspended device, while leaving the primary device in its current state. After you issue this command, the configuration is symmetrical with both devices being primaries. You can issue the *failbackpprc* command to resume mirroring in either direction. For the purposes of this demonstration, the *failbackpprc* command is issued to make Site A the primary again, thus overwriting any changes made at Site B.

When this command is issued, the following processing occurs:

1. The B volumes become suspended primary volumes.
2. Allow I/O to start at Site B.

Command example 30: Failoverpprc to allow the target to become read write enabled

DATA_DG_8K

```
failoverpprc -dev IBM.2107-75DF161 -remotedev IBM.2107-75DF162 -type mmir
7320:6220 7321:6221 7322:6222
```

LOGS_DG_8K

```
failoverpprc -dev IBM.2107-75DF161 -remotedev IBM.2107-75DF162 -type mmir
7530:6430 7531:6431 7532:6432
```

ARCH_DG_8K

```
failoverpprc -dev IBM.2107-75DF161 -remotedev IBM.2107-75DF162 -type mmir
7740:6640 7741:6641 7742:6642
```

Return response from *failoverpprc*:

```
dscli> failoverpprc -dev IBM.2107-75DF161 -remotedev IBM.2107-75DF162 -type
mmir 7320:6220 7321:6221 7322:6222
Date/Time: August 27, 2008 9:39:25 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF161
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7320:6220 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7321:6221 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7322:6222 successfully
reversed.
```

After issuing the *failoverpprc* commands, *lspprc* to view the status of the relationship on Site B

```
dscli> lspprc -l -fmt delim -delim | 7320-7322
Date/Time: August 27, 2008 9:41:50 AM PDT IBM DSCLI Version: 5.3.1.101 DS:
IBM.2107-75DF161
ID|State|Reason|Type|Out Of Sync Tracks|Tgt Read|Src Cascade|Tgt Cascade|Date
Suspended|SourceLSS|Timeout (secs)|Critical Mode|First Pass Status|
Incremental Resync|Tgt Write|GMIR CG|PPRC CG|AllowTgtSE|DisableAutoResync
==========================================================================
==========================================================================
==========================================================================
7320:6220|Suspended|Host Source|Metro Mirror|0|Disabled|Disabled|Invalid|-
|73|60|Disabled|Invalid|Disabled|Disabled|Disabled|Disabled|Disabled|-
7321:6221|Suspended|Host Source|Metro Mirror|0|Disabled|Disabled|Invalid|-
|73|60|Disabled|Invalid|Disabled|Disabled|Disabled|Disabled|Disabled|-
7322:6222|Suspended|Host Source|Metro Mirror|0|Disabled|Disabled|Invalid|-
|73|60|Disabled|Invalid|Disabled|Disabled|Disabled|Disabled|Disabled|-
```

The PPRC paths from Site A to Site B have been removed.

Reestablish Site A to Site B Metro Mirror

When activities on the remote site Site B are complete and you want to refresh the image of the data from the primary site Site A to the remote Site B. You need to perform the following steps:

1. Quiesce I/O at Site B.
2. Reestablish paths between the local and remote site LSSs that contain the Metro Mirror volume pairs.
3. Issue a *failbackpprc* on Site A to the A to B volume pairs.

Recreate pprcpaths (removed when freezepprc is issued)

Command example 31: Recreate pprcpaths after freezepprc

DATA_DG_8K

```
mkpprcpath -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -remotewwnn
5005076307FFC77D -consistgrp -srclss 62 -tgtlss 73 I0231:I0031
```

LOGS_DG_8K

```
mkpprcpath -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -remotewwnn
5005076307FFC77D -consistgrp -srclss 64 -tgtlss 75 I0231:I0031
```

ARCH_DG_8K

```
mkpprcpath -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -remotewwnn
5005076307FFC77D -consistgrp -srclss 66 -tgtlss 77 I0231:I0031
```

Failbackpprc Site A to Site B

At the local site Site A, issue a *failbackpprc* command for the A to B volume pairs.

When this command processes, the following occurs:

1. Updates that are made to the volumes at Site B are recorded with the change recording feature. Changed tracks of data are copied from the Site A volumes to Site B volumes.
2. When the copy process is complete, the Site B volumes are synchronized with the Site A volumes.

Command example 32: Restore Site A to Site B Metro Mirror relationship

```
failbackpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type mmir
6220:7320 6221:7321 6222:7322

failbackpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type mmir
6430:7530 6431:7531 6422:7532

failbackpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type mmir
6640:7720 6641:7741 6642:7742
```

Copy Services Backup and Recovery by example

This section describes the actual processes followed in the lab for the scenarios using the configurations and commands described in the previous sections.

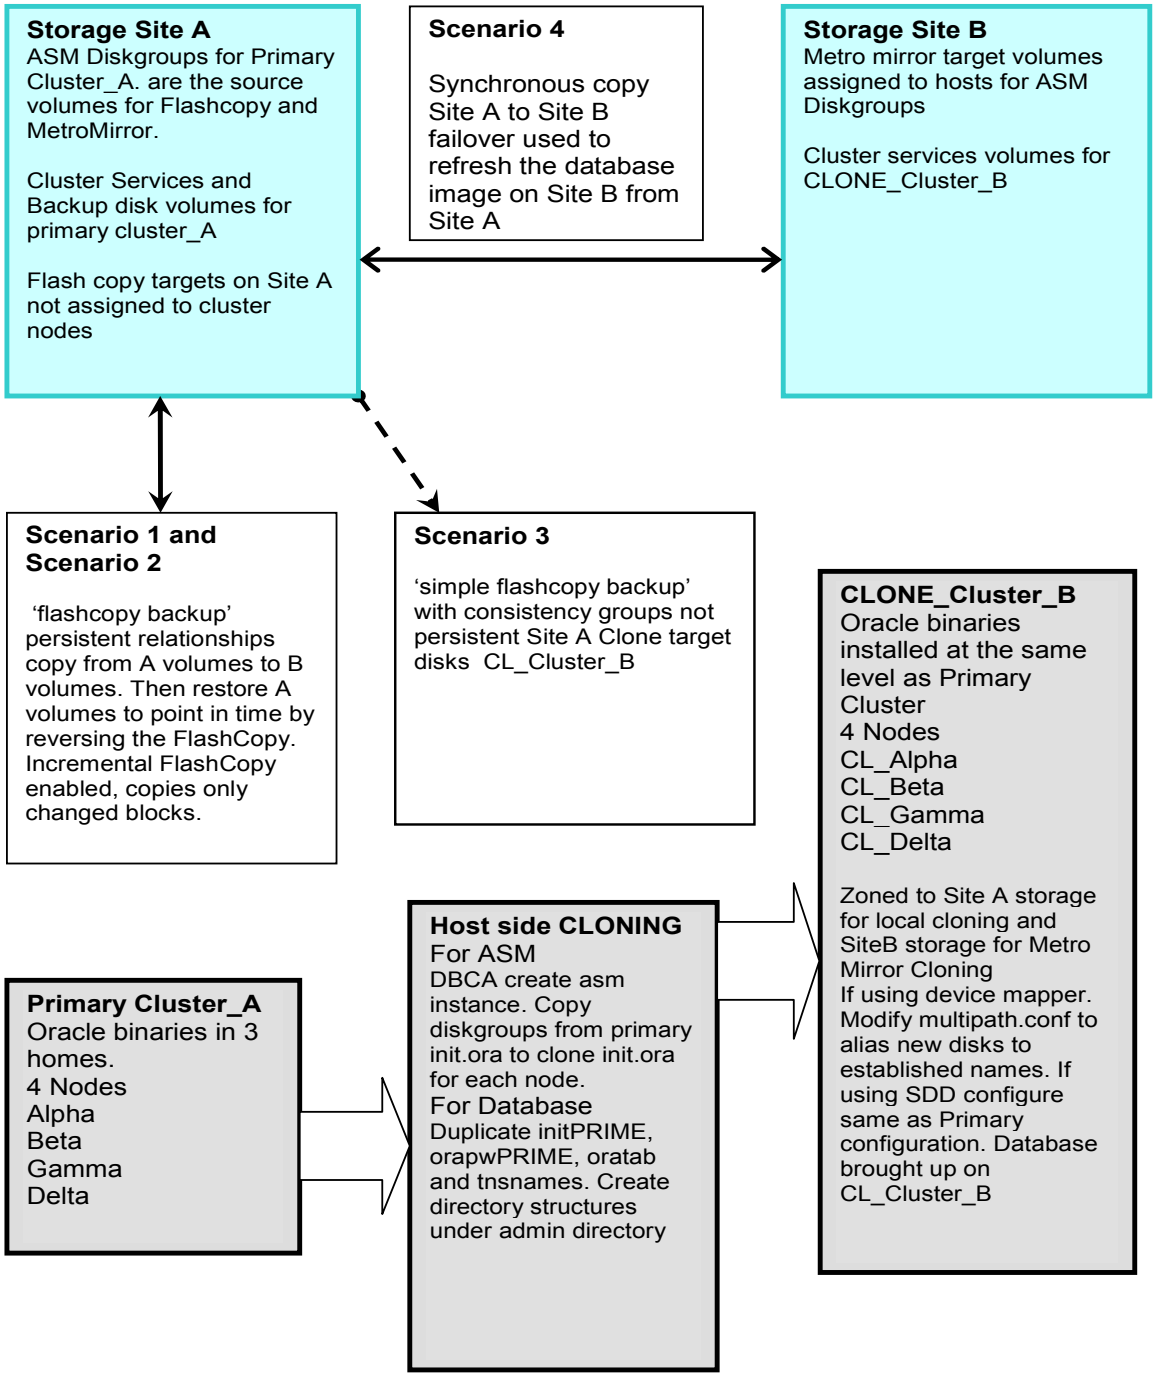

Figure 1: Lab processes for the scenarios

Backup and Recovery Procedures and Validations

The following scenarios are step-by-step descriptions of how to use DS8000 Advanced Copy Services to backup and recover Oracle RAC 10g databases on ASM. The usage of tape backups, which is a part of the backup strategy, is left up to the individual customer to incorporate into their own backup and recovery strategy.

The procedures describe in detail the steps used, integrating the Advanced Copy services processes with the Oracle backup and recovery process into one complete checklist.

The general approach that is taken is thus:

- Backup a running Oracle RAC 10g database on ASM using Advanced Copy Services.
- Restore the database using Advanced Copy Services.
- Recover the database by using standard Oracle database recovery procedures.
- Verify that the database is up and running and functioning normally.

For the scenarios, example commands are displayed below each step as follows:

DATA_DG_8K volume_id range 6220-6222 LSS 62

```
dscli>resyncflash -persist -record -freeze -seqnum 20 6220:3520 6221:3521
6222:3522
dscli>unfreezeflash 62
```

LOGS_DG_8K volume_id range 6430-6432 LSS 64

```
dscli>resyncflash -persist -record -freeze -seqnum 30 6430:3530 6431:3531
6432:3532
dscli>unfreezeflash 64
```

ARCH_DG_8K volume_id range 6640-6642 LSS 66

```
dscli>resyncflash -persist -record -freeze -seqnum 40 6640:3540 6641:3541
6642:3542
dscli>unfreezeflash 66
```

Scenario 1: Loss of DATA_DG_8K containing datafiles

In this scenario, database recovery is shown by reversing a FlashCopy mapping to restore a previous FlashCopy backup of the DATA_DG_8K ASM diskgroup. This assumes that the database can be recovered with the command “recover database” and that it can then be opened with “alter database open”. In other words, it assumes that the controlfiles, the online redo logs, and the spfile are not lost.

Prerequisites

Ensure that the following conditions are met before starting the tasks for this scenario:

- Database is configured and running (see **Table 4. DS8000 primary RAC database volume configuration**).
- DS8000 and ASM mappings are created (see **Table 8. ASM diskgroup configuration and contents**).
- FlashCopy mappings are established (see **Command example 13: Mkflash**).

The following is an outline of the steps that will be followed:

1. Backup ASM diskgroups DATA_DG_8K and LOGS_DG_8K.
2. Backup parameter and controlfiles.
3. Shut down database.
4. Reverse FlashCopy for restore.
5. Recover database on one node of the cluster.
6. After recovery is complete, start cluster resources for the database.

Recovering the ASM diskgroup DATA_DG_8K

1. Start a SQL*Plus session as SYSDBA and put the database in hot backup mode:

```
[oracle@ALPHA ~]$ export ORACLE_HOME=/u01/db/product/10.2/db
[oracle@ALPHA ~]$ export ORACLE_SID=PRIME8K/
[oracle@ALPHA ~]$ sqlplus / as sysdba
SQL>  alter database begin backup;
```

2. Resynchronize the FlashCopy for the DATA_DG_8K and LOGS_DG_8K consistency groups.

```
dscli> resyncflash -persist -record -freeze -seqnum 20 6220-6222:3520-3522
dscli> resyncflash -persist -record -freeze -seqnum 30 6430-6432:3530-3532
dscli> unfreezeflash 62 64
```

3. Take the database out of hot backup mode. Force the archiving of the online redo logs for all RAC threads.

```
SQL> alter database end backup;
SQL> alter system archive log current;
```

4. This database recovery will use the archived logs stored in the ASM diskgroup ARCH_DG_8K. The archived logs are backed up as a standard practice.

```
RMAN> backup archivelog all format'/u01/backup/PRIME8K/al_t%t_s%s_p%p';
```

5. Backup the controlfile.

```
RMAN>  sql 'alter database backup controlfile to trace';
RMAN>  backup as copy
       format'/u01/backup/PRIME8K/ctl_copy_t%t_s%s_p%p.ctl'
       current controlfile;
RMAN>  backup
       format '/u01/backup/PRIME8K/cf_t%t_s%s_p%p'
       current controlfile;
```

6. Backup the spfile.

```
RMAN>  sql "create pfile=''/u01/backup/PRIME8K/init.ora.back'' from
spfile";
RMAN>  backup
       format '/u01/backup/PRIME8K/spfile_t%t_s%s_p%p'
       spfile;
```

7. Verify the FlashCopy status ensuring OutOfSync Tracks=0 and Active Copy is Disabled. You are now ready to restore the diskgroup.

```
dscli>  lsflash -l -fmt delim -delim | -seqnum 20 6220-6222
```

8. Use srvctl on node ALPHA to shut down the database and all ASM instances:

```
[oracle@ALPHA ~]$  srvctl stop database -d PRIME8K -o abort
[oracle@ALPHA ~]$  srvctl stop asm -n  ALPHA -o abort
[oracle@ALPHA ~]$  srvctl stop asm -n  BETA -o abort
[oracle@ALPHA ~]$  srvctl stop asm -n  GAMMA -o abort
[oracle@ALPHA ~]$  srvctl stop asm -n  DELTA -o abort
```

9. The FlashCopy can be reversed only when the previous resyncflash command is complete, OutOfSyncTracks=0 and ActiveCopy is Disabled.

```
dscli> reverseflash -persist -record –tgtpprc -seqnum 22 6220-6222:3520-
3522
```

10. If using Linux: With asmlib on Linux, rescan the disks to refresh the disk information on all server nodes.

```
[root@ALPHA ~]$  /etc/init.d/oracleasm scandisks
[root@BETA  ~]$  /etc/init.d/oracleasm scandisks
[root@GAMMA ~]$  /etc/init.d/oracleasm scandisks
[root@DELTA ~]$  /etc/init.d/oracleasm scandisks
```

11. Start the ASM instance on node ALPHA.

```
[oracle@ALPHA ~]$  srvctl start asm -n ALPHA
```

12. Set the environment to the +ASM1 instance. Start SQL*Plus and check the status of the ASM diskgroups. Mount any diskgroups which are not already mounted:

```
[oracle@ALPHA ~]$  export ORACLE_HOME=u01/asm/product/10.2/asm
[oracle@ALPHA ~]$  export ORACLE_SID=ASM1
[oracle@ALPHA ~]$  sqlplus / as sysdba

SQL> select name, state from v$asm_diskgroup;
NAME       STATE
------------------------------ -----------
ARCH_DG_8K MOUNTED
DATA_DG_8K DISMOUNTED
LOGS_DG_8K MOUNTED
```

```
SQL> alter diskgroup DATA_DG_8K mount;
```

13. Startup mount the database instance on node ALPHA.

```
[oracle@ALPHA ~]$  srvctl start instance -i PRIME8K1 -d PRIME8K -o mount
```

14. Set environment to the PRIME8K1 instance. Start SQL*Plus and check the status of the datafiles. Bring any offline datafiles online:

```
[oracle@ALPHA ~]$  export ORACLE_HOME=u01/db/product/10.2/db
[oracle@ALPHA ~]$  export ORACLE_SID=PRIME8K1
[oracle@ALPHA ~]$  sqlplus / as sysdba
SQL> select name, status from v$datafile where status not in ('ONLINE',
'SYSTEM');
NAME  STATUS
------------------------------------------------- -------
+DATA_DG_8K/PRIME8K/datafile/soe.266.644618689  OFFLINE
SQL> alter database datafile
'+DATA_DG_8K/PRIME8K/datafile/soe.266.644618689' online;
```

15. Begin database recovery and specify automatic mode:

```
SQL> recover database;
ORA-00279: change 27391084 generated at 08/05/2008 08:47:07 needed for
thread 1
ORA-00289: suggestion : +ARCH_DG/PRIME8K/1_193_644615537.arc
ORA-00280: change 27391084 for thread 1 is in sequence #193
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
auto
```

16. When recovery is complete, open the database.

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

17. Use srvctl to start the remaining ASM instances, then start the remaining database instances.

```
[oracle@ALPHA ~]$  srvctl start asm -n  BETA
[oracle@ALPHA ~]$  srvctl start asm -n  GAMMA
[oracle@ALPHA ~]$  srvctl start asm -n  DELTA
[oracle@ALPHA ~]$  srvctl start database -d PRIME8K
```

18. Verify all cluster resources have returned to online using crs_stat -t

```
[root@ALPHA ~]$  crs_stat -t
```

19. Verify the FlashCopy status, ensuring OutOfSyncTracks=0 and ActiveCopy is Disabled for ASM diskgroup DATA_DG_8K volumes.

```
dscli> lsflash -l -seqnum 22 3520-3522
```

20. Once the database has been restored and recovered and ActiveCopy is Disabled for ASM diskgroup DATA_DG_8K volumes. Once the database has been restored and recovered. Modify the FlashCopy relationships back to the original relationship to prepare for normal FlashCopy backups.

```
dscli> reverseflash -persist -record -seqnum 20 3520-3522:6220-6222
```

21. Normal FlashCopy backups can now resume.

Scenario 2: Loss of the entire database

This exercise shows how to backup all ASM diskgroups in order to be able to recover from the loss of not only the database datafiles, but also the controlfiles, the online redo logs and the spfile. This is an incomplete recovery to a point-in-time before the loss of the ASM diskgroups.

Prerequisites

Ensure that the following conditions are met before starting the tasks for this scenario:

- The database is configured and running (see **Table 4. DS8000 primary RAC database volume configuration**).
- DS8000 and ASM mappings are created (see **Table 8. ASM diskgroup configuration and contents**).
- FlashCopy mappings are established (see **Command example 13: Mkflash).**

The following list is an outline of the steps that will be followed:

1. Backup ASM diskgroup LOGS_DG_8K.
2. Backup ASM diskgroup DATA_DG_8K.
3. Backup parameter and controlfiles.
4. Backup ASM diskgroup ARCH_DG_8K.
5. Shutdown the database.
6. Reverse the FlashCopy.
7. Recover the database on one node of the cluster.
8. After recovery is complete, start the remaining cluster resources for the database.

Recovering from the loss of all ASM diskgroups

1. Start a SQL*Plus session as SYSDBA and put the database in hot backup mode:

```
[oracle@ALPHA ~]$ export  ORACLE_HOME=/u01/db/product/10.2/db
[oracle@ALPHA ~]$ export ORACLE_SID=PRIME8K1
[oracle@ALPHA ~]$ sqlplus / as sysdba
SQL>    alter database begin backup;
```

2. Resync the FlashCopy relationships for the ASM diskgroups DATA_DG_8K and LOGS_DG_8K

```
dscli>   resyncflash -persist -record -freeze -seqnum 20 6220-6222:3520-
3522
dscli>   unfreezeflash 62
dscli>   resyncflash -persist -record -freeze -seqnum 30 6430-6432:3530-
3532
dscli>   unfreezeflash 64
```

3. Verify the FlashCopy status noting *OutOfSyncTracks* and *ActiveCopy.*

```
dscli>   lsflash -l -fmt delim -delim | 6220-6432
```

4. Take the database out of hot backup mode. Force the archiving of the online redo logs for all RAC threads.

```
SQL> alter database end backup;
SQL> alter system archive log current;
```

5. This database recovery will use the archived logs stored in the ASM diskgroup, ARCH_DG_8K.

```
RMAN>
 backup archivelog all
      format'/u01/backup/PRIME8K/al_t%t_s%s_p%p';
```

6. Backup the controlfile.

```
RMAN> sql 'alter database backup controlfile to trace';
RMAN> backup as copy
       format '/u01/backup/PRIME8K/ctl_copy_t%t_s%s_p%p.ctl'
       current controlfile;
RMAN> backup
       format '/u01/backup/PRIME8K/cf_t%t_s%s_p%p'
       current controlfile;
RMAN> backup
       format '+ARCH_DG_8K/PRIME8K/rman_backups/cf_t%t_s%s_p%p'
       current controlfile;
```

7. Backup the spfile.

```
RMAN>  sql "create pfile=''/u01/backup/PRIME8K/init.ora.back'' from
spfile";
RMAN>  backup
        format '/u01/backup/PRIME8K/spfile_t%t_s%s_p%p' spfile;
RMAN>  backup
        format '+ARCH_DG_8K/PRIME8K/rman_backups/spfile_t%t_s%s_p%p'
spfile;
```

8. Start the incremental FlashCopy for the ASM diskgroup ARCH_DG_8K volumes.

```
dscli>  resyncflash -persist -record -freeze -seqnum 40 6640-6642:3540-3542
dscli>  unfreezeflash 66
```

9. Prepare for a complete restore of all the ASM related volumes. Use srvctl on node ALPHA to shut down the database and all ASM instances.

```
[oracle@ALPHA ~]$ srvctl stop database -d PRIME8K -o abort
[oracle@ALPHA ~]$ srvctl stop asm -n ALPHA -o abort
[oracle@ALPHA ~]$ srvctl stop asm -n BETA -o  abort
[oracle@ALPHA ~]$ srvctl stop asm -n GAMMA -o abort
[oracle@ALPHA ~]$ srvctl stop asm -n DELTA -o abort
```

10. Monitor the the copy relationship and status information. Verify the FlashCopy status ensuring OutOfSyncTracks=0 and ActiveCopy is Disabled for all diskgroups.

```
dscli>  lsflash -l -fmt delim -delim | 6220-6642
```

11. Reverse the FlashCopy relationships to perform the restore for the DATA_DG_8K, ARCH_DG_8K, and LOGS_DG_8K volumes.

```
dscli>  reverseflash -persist -record –tgtpprc -seqnum 22 6220-6222:3520-
3522
dscli>  reverseflash -persist -record –tgtpprc -seqnum 33 6430-6432:3530-
3532
dscli>  reverseflash -persist -record –tgtpprc -seqnum 44 6640-6642:3540-
3542
```

12. Monitor the the copy relationship and status information.

```
dscli>  lsflash -l -fmt delim -delim | 3520-3542
```

13. If using Linux: With asmlib on Linux, rescan the disks to refresh the disk information on all server nodes after the reverseflash has been issued.

```
[root@ALPHA ~]$ /etc/init.d/oracleasm scandisks
[root@BETA  ~]$ /etc/init.d/oracleasm scandisks
[root@GAMMA ~]$ /etc/init.d/oracleasm scandisks
[root@DELTA ~]$ /etc/init.d/oracleasm scandisks
```

14. Start the ASM instance on node ALPHA.

```
[oracle@ALPHA ~]$ srvctl start asm -n ALPHA
```

15. Set the environment to the +ASM1 instance. Start SQL*Plus and check the status of the ASM diskgroups. Mount any diskgroups which are not already mounted:

```
[oracle@ALPHA ~]$ export ORACLE_HOME=/u01/asm/product/10.2/asm
[oracle@ALPHA ~]$ export ORACLE_SID=ASM1
[oracle@ALPHA ~]$ sqlplus / as sysdba
SQL> select name, state from v$asm_diskgroup;
NAME                           STATE
------------------------------ -----------
DATA_DG_8K                     DISMOUNTED
LOGS_DG_8K                     MOUNTED
ARCH_DG_8K                     MOUNTED

SQL> alter diskgroup DATA_DG_8K mount;
```

16. Startup nomount the database instance on node ALPHA.

```
[oracle@ALPHA ~]$ srvctl start instance -i PRIME8K1 -d PRIME8K -o nomount
```

17. Restore the controlfile.

```
RMAN> restore controlfile;
```

18. Set the environment to the PRIME8K1 instance. Then start SQL*Plus, mount the database and check the status of the datafiles. Bring any offline datafiles online.

```
[oracle@ALPHA ~]$ export ORACLE_HOME=u01/db/product/10.2/db
[oracle@ALPHA ~]$ export ORACLE_SID=PRIME8K1
[oracle@ALPHA ~]$ sqlplus / as sysdba
SQL> alter database mount;
SQL> select name, status from v$datafile where status not in ('ONLINE',
'SYSTEM');
NAME                                               STATUS
-------------------------------------------------  ----------
+DATA_DG_8K/prime8k/datafile/soe.266.644618689     OFFLINE
SQL> alter database datafile
‘+DATA_DG_8K/prime8k/datafile/soe.266.644618689’ online;
```

19. Begin the database recovery and specify *automatic mode* on node ALPHA.

```
SQL> recover database until cancel using backup controlfile;
ORA-00279: change 27391084 generated at 01/25/2008 08:47:07 needed for
thread 1
ORA-00289: suggestion : +ARCH_DG_8K/prime8k/1_193_644615537.arc
ORA-00280: change 27391084 for thread 1 is in sequence #193
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
auto
```

20. When recovery is complete, open the database resetting the logs.

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

21. Use Oracle Clusterware to start the remaining ASM instances, then start the remaining database instances.

```
[oracle@ALPHA ~]$ srvctl start asm -n  BETA
[oracle@ALPHA ~]$ srvctl start asm -n  GAMMA
[oracle@ALPHA ~]$ srvctl start asm -n  DELTA
[oracle@ALPHA ~]$ srvctl start database -d PRIME8K
```

22. Verify all cluster services have returned to online.

```
[root@ALPHA ~]$ crs_stat -t
```

23. Verify the FlashCopy status, ensuring OutOfSyncTracks=0 and ActiveCopy is Disabled for all ASM diskgroup volumes.

```
dscli>  lsflash -l -fmt delim -delim | 3520-3542
```

24. Modify the FlashCopy relationships back to the original relationship.

```
dscli>  reverseflash -persist -record -seqnum 20 3520-3522:6220-6222
dscli>  reverseflash -persist -record -seqnum 30 3530-3532:6430-6432
dscli>  reverseflash -persist -record -seqnum 40 3540-3542:6640-6642
```

25. Normal FlashCopy backups can now resume.

Scenario 3: FlashCopy Clone of a Oracle RAC 10g Database with ASM

This exercise shows how to make a copy (clone) of an Oracle RAC 10g database with ASM to another Oracle RAC cluster. It is expected the target RAC nodes have been prepared with the same Oracle binaries installed on the source RAC cluster. This example does not rename the database, this cloning process could be used to set up a database for testing or development and the additional process of purposing the copied database is not covered.

Prerequisites

- The source database is configured and running (see **Table 4. DS8000 primary RAC database volume configuration)**.
- The target RAC cluster is set up, (see **Table 9. Clone and Remote Oracle RAC node configuration** and **Table 5. DS8000 Clone and Remote RAC volume configuration**)
- The LUNs are configured and visible on the target RAC cluster, (see **Table 8 ASM diskgroup configuration for Clone Remote RAC cluster**).
- The ASM setup is complete, (see **Clone or Remote Oracle ASM setup)**
- FlashCopy mappings are prepared (see **Table 12 Clone FlashCopy mappings)**

The following list is an outline of the steps that will be followed:

1. Trigger FlashCopy consistency groups to create a consistent backup.
2. Register the resources for the copied database with the target CRS.
3. Recover the database on one node of the target cluster.
4. After recovery is complete, start the remaining cluster resources for the database.

Cloning Oracle RAC 10g with ASM

1. Start a SQL*Plus session as SYSDBA and put the database in hot backup mode:

```
[oracle@ALPHA ~]$ export  ORACLE_HOME=/u01/db/product/10.2/db
[oracle@ALPHA ~]$ export ORACLE_SID=PRIME8K1
[oracle@ALPHA ~]$ sqlplus / as sysdba
SQL> alter database begin backup;
```

2. Start the incremental FlashCopy for the DATA_DG_8K and LOGS_DG_8K consistency groups.

```
dscli>  mkflash -freeze -seqnum 20 6220-6222:7320-7322
dscli>  mkflash -freeze -seqnum 30 6430-6432:7530-7532
dscli>  unfreezeflash 62 64
```

3. Check the copy relationships and status information.

```
dscli>  lsflash -l -fmt delim -delim | 6220-6432
```

4. Take the database out of hot backup mode. Force archival of the online redo logs for all RAC threads.

```
SQL>   alter database end backup;
SQL>   alter system archive log current;
```

5. Monitor the copy relationship and status information.

```
dscli>  lsflash -l -fmt delim -delim | 6220-6432
```

6. Backup the archived logs to the mounted backup directory.

```
RMAN>  backup archivelog all
       format'/u01/backup/PRIME8K/al_t%t_s%s_p%p';
```

7. Backup the controlfile.

```
RMAN>  sql 'alter database backup controlfile to trace';

RMAN>  backup as copy
       format '/u01/backup/PRIME8K/ctl_copy_t%t_s%s_p%p.ctl'
       current controlfile;

RMAN>  backup
       format '/u01/backup/PRIME8K/cf_t%t_s%s_p%p'
       current controlfile;

RMAN>  backup
       format '+ARCH_DG_8K/PRIME8K/rman_backups/cf_t%t_s%s_p%p'
       current controlfile;
```

8. Backup the spfile.

```
RMAN> sql "create pfile=''/u01/backup/PRIME8K/init.ora.back'' from spfile";
RMAN> backup
      format '/u01/backup/PRIME8K/spfile_t%t_s%s_p%p'
      spfile;
RMAN> backup
      format '+ARCH_DG_8K/PRIME8K/rman_backups/spfile_t%t_s%s_p%p'
      spfile;
```

9. Start the incremental FlashCopy for the ARCH_DG_8K consistency group.

```
dscli>  mkflash -freeze -seqnum 40 6640-6642:7740-7742
dscli>  unfreezeflash 66
```

10. Check the status of the copy, note the OutOfSyncTracks and ActiveCopy status information.

```
dscli>   lsflash -l -fmt delim -delim | 6220-6642
```

11. With ASMLib on Linux, rescan the disks to refresh the disk information on all Clone RAC nodes.

```
[root@CL_ALPHA ~]$ /etc/init.d/oracleasm scandisks
[root@CL_BETA  ~]$ /etc/init.d/oracleasm scandisks
[root@CL_GAMMA ~]$ /etc/init.d/oracleasm scandisks
[root@CL_DELTA ~]$ /etc/init.d/oracleasm scandisks
```

12. Start the ASM instance on the Clone node CL_ALPHA.

```
[oracle@CL_ALPHA ~]$ srvctl start asm -n CL_ALPHA
```

13. Set the environment to the +ASM1 instance on the clone node, CL_ALPHA. Start SQL*Plus and check the status of the ASM diskgroups. Mount any diskgroups not already mounted:

```
[oracle@CL_ALPHA ~]$ export ORACLE_HOME=u01/asm/product/10.2/asm
[oracle@CL_ALPHA ~]$ export ORACLE_SID=ASM1
[oracle@CL_ALPHA ~]$ sqlplus / as sysdba
SQL> select name, state from v$asm_diskgroup;
NAME                           STATE
------------------------------ -----------
DATA_DG_8K                     DISMOUNTED
LOGS_DG_8K                     MOUNTED
ARCH_DG_8K                     MOUNTED
```

```
SQL> alter diskgroup DATA_DG_8K mount;
```

14. On the clone node CL_ALPHA, add the database to the cluster registry on the clone cluster.

```
[oracle@CL_ALPHA ~]$ export ORACLE_HOME=/u01/db/product/10.2/db
[oracle@CL_ALPHA ~]$ srvctl add database -d PRIME8K -o
/u01/db/product/10.2/db
```

15. On the clone node, CL_ALPHA add the instances to the cluster registry on the clone cluster.

```
[oracle@CL_ALPHA ~]$ srvctl add instance -d PRIME8K -i PRIME8K1 -n CL_ALPHA
[oracle@CL_ALPHA ~]$ srvctl add instance -d PRIME8K -i PRIME8K2 -n CL_BETA
[oracle@CL_ALPHA ~]$ srvctl add instance -d PRIME8K -i PRIME8K3 -n CL_GAMMA
[oracle@CL_ALPHA ~]$ srvctl add instance -d PRIME8K -i PRIME8K4 -n CL_DELTA
```

16. Start the database instance on the Clone node, CL_ALPHA using the nomount option.

```
[oracle@CL_ALPHA ~]$ srvctl start instance -i PRIME8K1 -d PRIME8K -o
nomount
```

17. Restore the controlfile on the Clone node, CL_ALPHA.

```
RMAN> restore controlfile;
```

18. Set the environment to the PRIME8K1 instance on the Clone node CL_ALPHA. Start SQL*Plus, mount the database and check the status of the datafiles. Bring any offline datafiles online.

```
[oracle@CL_ALPHA ~]$ export ORACLE_HOME=u01/db/product/10.2/db
[oracle@CL_ALPHA ~]$ export ORACLE_SID=PRIME8K1
[oracle@CL_ALPHA ~]$ sqlplus / as sysdba

SQL> alter database mount;

SQL> select name, status from v$datafile where status not in ('ONLINE',
'SYSTEM');

NAME                                               STATUS
-------------------------------------------------  ----------
+DATA_DG_8K/prime8k/datafile/soe.266.644618689     OFFLINE

SQL> alter database datafile
'+DATA_DG_8K/prime8k/datafile/soe.266.644618689' online;
```

19. Begin the database recovery and specify automatic mode on the Clone node CL_ALPHA.

```
SQL> recover database until cancel using backup controlfile;

ORA-00279: change 27391084 generated at 01/25/2008 08:47:07 needed for
thread 1
ORA-00289: suggestion : +ARCH_DG_8K/prime8k/1_193_644615537.arc
ORA-00280: change 27391084 for thread 1 is in sequence #193
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}

auto
```

20. When recovery is complete, use 'alter database open resetlogs' on Clone node CL_ALPHA.

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

21. Use srvctl to start the remaining ASM instances and then start the remaining database instances on the Clone RAC nodes.

```
[oracle@CL_ALPHA ~]$ srvctl start asm -n  BETA
[oracle@CL_ALPHA ~]$ srvctl start asm -n  GAMMA
[oracle@CL_ALPHA ~]$ srvctl start asm -n  DELTA
[oracle@CL_ALPHA ~]$ srvctl start database -d PRIME8K
```

22. Verify all Clone cluster services have returned to online using crs_stat -t.

```
[root@CL_ALPHA ~]$ crs_stat -t
```

Scenario 4: Metro Mirror Remote Clone of a Oracle RAC 10g database

This exercise shows how to clone a Oracle RAC 10g database on ASM to a remote DS8000 and Oracle RAC cluster using Metro Mirror.

Prerequisites

Ensure that the following conditions are met before starting the database recovery tasks for this scenario:

- The PPRC paths are established (see **Creating storage unit relationships**).
- The source database is configured and running (see **Table 4. DS8000 primary RAC database volume configuration**).
- The remote RAC cluster is set up, (see **Table 9 Clone and remote Oracle RAC node configuration** and **Table 5. DS8000 Clone and Remote RAC volume configuration**).
- The LUNs are configured and visible on the remote RAC cluster, (see **Table 10. ASM diskgroup configuration for Clone Oracle RAC cluster**).
- The ASM setup is complete, (see **Clone or Remote Oracle ASM setup**)
- The database instance configuration is completed, **(**see **Clone or Remote Database Instance setup**).
- The Metro Mirror relationships are created, (see **Metro Mirror configuration**).

The following list is an outline of the steps that will be followed:

1. Create Metro Mirror relationships with the *-consistency* parameter.
2. Freeze the LSS for the Metro Mirror consistency groups when the state is full-duplex.
3. Register the resources for the copied database with the remote CRS.
4. Recover the database on one node of the remote Oracle RAC cluster.
5. After recovery is complete, start the remaining cluster resources for the database.

Remote Cloning of Oracle RAC 10g with ASM

1. Verify the status of the Metro Mirror relationships:

   DATA Source LSS 62 relationships 6220:7320 6221:7321 6222:7322

   LOGS Source LSS 64 relationships 6430:7530 6431:7531 6432:7532

   ARCH Source LSS 66 relationships 6640:7740 6641:7741 6642:7742

```
dscli> lspprc -l -fmt delim -delim | 6220-6642
```

2. When the status of all Metro Mirror consistency groups shows full duplex, you are ready to begin the process in this scenario.
3. On the source RAC cluster node, start a SQL*Plus session as SYSDBA and put the database in hot backup mode.

```
[oracle@ALPHA ~]$ export $ORACLE_HOME=/u01/db/product/10.2/db
[oracle@ALPHA ~]$ export $ORACLE-SID=PRIME
[oracle@ALPHA ~]$ sqlplus / as sysdba
SQL> alter database begin backup;
```

4. Freeze the DATA_DG_8K Metro Mirror consistency group. Then unfreeze the LSS.

```
dscli> freezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 62:73
dscli> unfreezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 62:73
```

5. From the source RAC cluster, take the database out of backup mode and force the archiving of the online redo logs

```
SQL> alter database end backup;
SQL> alter system archive log current;
```

6. Freeze the LOGS_DG_8K consistency group. Then unfreeze the LSS.

```
dscli> freezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 64:75
dscli> unfreezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 64:75
```

7. As a standard practice, backup the archived logs to shared storage.

```
 RMAN>backup archivelog all format'/u01/backup/PRIME8K/al_t%t_s%s_p%p';
```

8. Backup the controlfile.

```
RMAN>  sql 'alter database backup controlfile to trace';
RMAN> backup
      format '/u01/backup/PRIME8K/ctl_copy_t%t_s%s_p%p.ctl'
      current controlfile;
RMAN> backup
      format '+ARCH_DG_8k/PRIME8K/rman_backups/cf_t%t_s%s_p%p'
      current controlfile;
```

9. Backup the spfile.

```
RMAN> sql "create pfile=''/u01/backup/PRIME8K/init.ora.back'' from spfile";
RMAN> backup
      format '+ARCH_DG_8k/PRIME8K/rman_backups/spfile_t%t_s%s_p%p'
      spfile;
```

10. Freeze the ARCH_DG_8K Metro Mirror consistency group.

```
dscli> freezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 66:77
dscli> unfreezepprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 66:77
```

11. Check the status of the Metro Mirror consistency group mappings on Site A. They should all show a status of suspended Freeze.

```
dscli> lspprc -l -fmt delim -delim | 6220-6642
Date/Time: June 24, 2008 5:15:38 PM PDT IBM DSCLI Version: 5.0.6.229 DS:
IBM.2107-75DF162
ID|State|Reason|Type|Out Of Sync Tracks|Tgt Read|Src Cascade|Tgt
Cascade|Date Suspended|SourceLSS|Timeout (secs)|Critical Mode|First Pass
Status
=============================================================================
=============================================================
6220:7320|Suspended|Freeze|Metro Mirror|361|Disabled|Disabled|invalid|-
|62|300|Disabled|Invalid
```

12. Execute the *failoverpprc* command enabling Site B I/O

```
dscli> failoverpprc -dev IBM.2107-75DF161  -remotedev IBM.2107-75DF162 -
type mmir 7320:6220 7321:6221 7322:6222

dscli> failoverpprc -dev IBM.2107-75DF161  -remotedev IBM.2107-75DF162 -
type mmir 7530:6430 7531:6431 7532:6432

dscli> failoverpprc -dev IBM.2107-75DF161  -remotedev IBM.2107-75DF162 -
type mmir 7740:6640 7741:6641 7742:6642
```

13. With ASMLib on Linux, run *oracleasm scandisks* to refresh the disk information on all clone RAC nodes.

```
[root@CL_ALPHA ~]$ /etc/init.d/oracleasm scandisks
[root@CL_BETA ~]$ /etc/init.d/oracleasm scandisks
[root@CL_GAMMA ~]$ /etc/init.d/oracleasm scandisks
[root@CL_DELTA ~]$ /etc/init.d/oracleasm scandisks
```

14. Start the ASM instance on Clone RAC node1.

```
[oracle@CL_ALPHA ~]$  srvctl start asm -n CL_ALPHA
```

15. As the user oracle on node CL_ALPHA, check the status of the ASM diskgroups. Mount any diskgroups which are not already mounted.

```
[oracle@CL_ALPHA ~]$  export ORACLE_HOME=u01/asm/product/10.2/asm
[oracle@CL_ALPHA ~]$  export ORACLE_SID=+ASM1
[oracle@CL_ALPHA ~]$  sqlplus / as sysdba
SQL> select name, state from v$asm_diskgroup;
NAME                           STATE
------------------------------ -----------
DATA_DG                        DISMOUNTED
LOGS_DG                        MOUNTED
ARCH_DG                        MOUNTED
SQL> alter diskgroup DATA_DG mount;
```

16. Add the database to the cluster registry on the Clone cluster.

```
[oracle@CL_ALPHA ~]$  srvctl add database -d PRIME8K -o
/u01/db/product/10.2/db
```

17. Add the instances to the database on the clone cluster.

```
[oracle@CL_ALPHA ~]$ srvctl add instance -d PRIME8K -i PRIME8K1 -n CL_ALPHA
[oracle@CL_ALPHA ~]$ srvctl add instance -d PRIME8K -i PRIME8K2 -n CL_BETA
[oracle@CL_ALPHA ~]$ srvctl add instance -d PRIME8K -i PRIME8K3 -n CL_GAMMA
[oracle@CL_ALPHA ~]$ srvctl add instance -d PRIME8K -i PRIME8K4 -n CL_DELTA
```

18. Startup nomount the database instance on cluster node CL_ALPHA.

```
[oracle@CL_ALPHA ~]$  srvctl start instance -i PRIME8K1 -d PRIME8K -o
nomount
```

19. On node CL_ALPHA, restore the controlfile. (RMAN will use the backup directory on ASM diskgroup +ARCH_DG.)

```
RMAN> restore controlfile;
```

20. On clone RAC node CL_ALPHA, mount the database.

```
[oracle@CL_ALPHA ~]$  export ORACLE_HOME=u01/db/product/10.2/db
[oracle@CL_ALPHA ~]$  export ORACLE_SID=PRIME8K1
[oracle@CL_ALPHA ~]$  sqlplus / as sysdba
```

```
SQL> alter database mount;
```

21. Check the status of the datafiles. Bring any offline datafiles online.

```
SQL> select name, status from v$datafile where status not in ('ONLINE',
'SYSTEM');
NAME                                               STATUS
-------------------------------------------------- -------
+DATA_DG/prime/datafile/soe.266.644618689          OFFLINE
SQL> alter database datafile '+DATA_DG/prime8k/datafile/soe.266.644618689'
online;
```

22. Begin database recovery and specify automatic mode.

```
SQL> recover database until cancel using backup controlfile;
ORA-00279: change 27391084 generated at 01/25/2008 08:47:07 needed for
thread 1
ORA-00289: suggestion : +ARCH_DG/prime8k/1_193_644615537.arc
ORA-00280: change 27391084 for thread 1 is in sequence #193

Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
auto
```

23. When recovery is complete, execute the *alter database open resetlogs* command on the database on clone node CL_ALPHA.

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

24. Use Oracle Clusterware to start the remaining ASM instances on the clone nodes, and then start the remaining database instances on the clone nodes.

```
[oracle@CL_ALPHA ~]$  srvctl start asm -n  CL_BETA
[oracle@CL_ALPHA ~]$  srvctl start asm -n  CL_GAMMA
[oracle@CL_ALPHA ~]$  srvctl start asm -n  CL_DELTA
[oracle@CL_ALPHA ~]$  srvctl start database -d PRIME8K
```

25. Verify all cluster resources have returned to online using crs_stat -t

```
[root@CL_ALPHA ~]$  crs_stat -t
```

26. When it is time to refresh the data on storage Site B from storage Site A shut down the database and stop the ASM instances on the remote RAC cluster.

```
[oracle@CL_ALPHA ~]$  srvctl stop database -d PRIME8K -o abort
[oracle@CL_ALPHA ~]$  srvctl stop asm -n CL_ALPHA -o abort
[oracle@CL_ALPHA ~]$  srvctl stop asm -n CL_BETA -o abort
[oracle@CL_ALPHA ~]$  srvctl stop asm -n CL_GAMMA -o abort
[oracle@CL_ALPHA ~]$  srvctl stop asm -n CL_DELTA -o abort
```

27. Recreate the PPRC paths that were removed when the *freeze* was issued.

```
dscli> mkpprcpath  -dev IBM.2107-75DF162  -remotedev IBM.2107-75DF161 -
remotewwnn 5005076307FFC2C8 -consistgrp -srclss 62  -tgtlss 73  I0230:I0103

dscli> mkpprcpath  -dev IBM.2107-75DF162  -remotedev IBM.2107-75DF161 -
remotewwnn 5005076307FFC2C8 -consistgrp -srclss 64  -tgtlss 75  I0230:I0103

dscli> mkpprcpath  -dev IBM.2107-75DF162  -remotedev IBM.2107-75DF161 -
remotewwnn 5005076307FFC2C8 -consistgrp -srclss 66  -tgtlss 77  I0230:I0103
```

28. Failback the Metro Mirror relationships

```
dscli> failbackpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type
mmir 6220:7320 6221:7321 6222:7322

dscli> failbackpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type
mmir 6430:7530 6431:7531 6432:7532

dscli> failbackpprc -dev IBM.2107-75DF162 -remotedev IBM.2107-75DF161 -type
mmir 6640:7740 6641:7741 6642:7742
```

Scenario 5: Backup and Restore the OCR and Voting disks

This is to test the situation where the Oracle Cluster Registry and voting disks are lost. FlashCopy is used to backup and restore the voting disks. The OCR cannot be backed up with FlashCopy. The OCR backup and recovery is handled by Oracle Clusterware. Oracle Clusterware will be running properly after the recovery.

Prerequisites for cluster registry backup and restore

Ensure that the following conditions are met before starting the recovery tasks for this scenario:

- Automatic cluster registry backups are working properly
- OCR exports are taken on a regularly scheduled basis by Oracle Clusterware
- FlashCopy relationships for the voting disks are set up as found in the following table

SOURCE LUN NAME	TARGET LUN NAME
OCR_6003	OCR_3503
OCR_6004	OCR_3504
OCR_6005	OCR_3505

The following is an outline of the steps that will be followed:

1. Verify that the automatic OCR backups are taking place
2. Run explicit backups of the OCR
3. Run regularly scheduled FlashCopy backups of the voting disks
4. Stop the autostart of Oracle Clusterware and reboot all nodes
5. Recover the OCR by using one of the automatic backups or by using one of the explicitly taken backups
6. Restore the FlashCopy backups of the voting disks
7. Restart Oracle Clusterware

Backup and restore the OCR and Vote disks

1. On a regular basis, display the location, timestamp, and the originating node name of the backup files that Oracle created in the past 4 hours, 8 hours, 12 hours, and in the last day and week. This output should be saved since the location of the backups will not be obtainable from the OCR disks if they are lost.

```
[root@ALPHA ~]$ ocrconfig -showbackup
```

2. Take an explicit export of the OCR on a regularly scheduled basis.

```
[root@ALPHA ~]$ ocrconfig -export ocr_export_$(date +'%Y_%m_%d_%H_%M')
```

3. Start the FlashCopy with the -freeze parameter.

```
dscli> mkflash -freeze 6003-6005:3503-3505
dscli> unfreezeflash 60
```

4. When it is time to start recovery of the OCR and Vote disks, disable the starting of Oracle Clusterware during system reboot.

```
[root@ALPHA ~]$ /etc/init.d/init.crs disable
[root@BETA ~]$ /etc/init.d/init.crs disable
[root@GAMMA ~]$ /etc/init.d/init.crs disable
[root@DELTA ~]$ /etc/init.d/init.crs disable
```

5. Reboot all Oracle RAC nodes.

```
[root@ALPHA ~]$ shutdown now -r
[root@BETA  ~]$ shutdown now -r
[root@GAMMA ~]$ shutdown now -r
[root@DELTA ~]$ shutdown now -r
```

6. Restore the OCR by using either one of the automatically taken backups or one of the explicitly taken “exports”.

    a. If restoring one of the automatic backups (recommended):

```
[root@ALPHA ~]$  ocrconfig -restore
/u01/crs/product/10.2/crs/cdata/Prime8k_crs/backup00.ocr
```

    b. If restoring one of the explicit exports:

```
[root@ALPHA ~]$  ocrconfig -import ocr_export_2008_08_30_13_25
```

7. Verify that the OCR disks are now available.

```
[root@ALPHA ~]$ ocrcheck
```

8. Recover the voting disks by starting the reverse FlashCopy mappings.

```
dscli> mkflash -freeze 3503-3505: 6003-6005
```

9. Verify that the voting disks are now available.

```
[root@ALPHA ~]$ crsctl query css votedisk
```

10. Re-enable the autostart of Cluster Services on all nodes.

```
[root@ALPHA ~]$ /etc/init.d/init.crs enable
[root@BETA ~]$ /etc/init.d/init.crs enable
[root@GAMMA ~]$ /etc/init.d/init.crs enable
[root@DELTA ~]$ /etc/init.d/init.crs enable
```

11. Start Cluster Services on all nodes.

```
[root@ALPHA ~]$ /etc/init.d/init.crs start
[root@BETA ~]$ /etc/init.d/init.crs start
[root@GAMMA ~]$ /etc/init.d/init.crs start
[root@DELTA ~]$ /etc/init.d/init.crs start
```

Summary

This white paper shows methods for using IBM System Storage DS8000 Advanced Copy Services to backup and recover an Oracle Real Applications Clusters 10g database using Automatic Storage Management.

The Technology section gives a conceptual overview of fundamental System Storage DS8000, Advanced Copy Services and Oracle RAC database (and its features), which you need to understand before implementing this strategy.

The Configuration and Setup section provides a detailed description of both the hardware and software infrastructure required to implement this architecture. Storage configuration is provided in a table format along with reasons provided at each step of the way of how best to accomplish User-Managed Backup and Recovery through the use of Advanced Copy Services, in conjunction with Oracle RMAN utility.

The Copy Services Backup and Recovery by Example section provides a step-by-step application of the integrated processes required to backup and recover an Oracle RAC database using ASM with Advanced Copy Services functions.

Overall, in each of the scenarios, one through four, the database is prepared for backup, the backup is performed, and then, a restore and recovery process is demonstrated.

Scenario 1 and Scenario 2 demonstrate using FlashCopy for backup and restore of the primary database on the same Oracle RAC and same storage.

Scenario 3 demonstrates the FlashCopy cloning of the database to a different RAC on the same storage unit as the primary database.

Scenario 4 demonstrates the usage of Metro Mirror to synchronously clone the database to a remote storage unit and bring it up on a remote RAC.

Finally, Scenario 5 demonstrates the often overlooked need to backup the Oracle Cluster Service and Vote disks.

Appendix A: Index of command examples and tables

Command examples

Tables

Appendix B: References

References to supplement the information contained in this document are provided here:

IBM Redbooks®

- www.redbooks.ibm.com

System Storage DS8000

- IBM System storage DS8000 Series: Architecture and Implementation
  www.redbooks.ibm.com/redbooks/pdfs/sg246786.pdf
- IBM System Storage DS8000: Copy Services in Open Environments
  www.redbooks.ibm.com/redpieces/pdfs/sg246788.pdf
- The IBM System Storage DS8000 Information Center
  http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp
- IBM System Storage Command Line Interface User's Guide
  www-1.ibm.com/support/docview.wss?uid=ssg1S7001162&loc=en_US&cs=utf-8&lang=en
- DS8000 Terms and Concepts
  www.redbooks.ibm.com/abstracts/tips0535.html?Open

Oracle

- Backup and Recovery Generic Guidelines with Point-in-Time Copy Solutions and Automatic Storage Management
  http://ioug.itconvergence.com/pls/apex/CILOUG.download_my_file?p_file=403
- Oracle Database Concepts 10g Release 2 (10.2) Part Number B14220-02
  http://download.oracle.com/docs/cd/B19306_01/server.102/b14220/toc.htm
- Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide 10g Release 2 (10.2) Part Number B14197-07
  http://download.oracle.com/docs/cd/B19306_01/rac.102/b14197/toc.htm
- Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide 10g Release 2 (10.2) Part No. B14201-04
  http://download.oracle.com/docs/cd/B19306_01/install.102/b14203/toc.htm
- Oracle Database Backup and Recovery Advanced User's Guide 10g Release 2 (10.2) Part No. B14191-02
  http://download.oracle.com/docs/cd/B19306_01/backup.102/b14191/toc.htm
- Oracle Database 10g Release 2 Automatic Storage Management Overview and Technical Best Practices
  www.oracle.com/technology/products/database/asm/pdf/asm_10gr2_bestpractices%2005-07.pdf

Trademarks and special notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.