

**Using Oracle Database 10g™
Automatic Storage Management
with NEC Storage S series products**

An Oracle and NEC White Paper

August 2005

Updated on April 2006

ORACLE®

NEC Empowered by Innovation

Copyright 2005 NEC and Oracle Corporation. All rights reserved.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT ARE DISCLAIMED.

Oracle, Oracle Database 10g and other product names of Oracle are trademarks, registered trademarks of Oracle Corporation in the U.S. and other countries.

Linux is a trademark and/or registered trademark of Linus Torvalds in the United States and other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Rev. 1.00 August 2005

Rev. 2.00 December 2005

Contents

1. Introduction	4
2. NEC Storage features	4
2.1 Dual Parity (RAID-6)	4
2.2 Dynamic Pool	5
2.3 NEC Path Manager	6
2.4 DynamicDataReplication	7
3. Oracle features	7
3.1 Automatic Storage Management	7
4. ASM configuration	9
4.1 Special File names for NEC Path Manager	9
4.2 RAW Device Setting for shared RAC devices	9
4.3 Creating ASM Disk Groups with External Redundancy	11
5. Best Practices for ASM with NEC Storage	11
5.1 Disk Groups with Dual parity (RAID6)	11
5.2 Use Path Manager with ASM for Improved HA and Performance	12
5.3 Use More Than One LUN per ASM Disk Group for Performance	13
5.4 Dynamic storage capacity expansion	13
6. Conclusion	15
7. Appendix	15
7.1 Resizing Disks within an ASM Disk Group	15
7.2 Backup and Recovery Using DynamicDataReplication (DDR)	18

1. Introduction

The purpose of this joint white paper is to provide best practices for use of the new Oracle Database 10g feature Automatic Storage Management (ASM) with NEC Storage S series products.

Automatic Storage Management (ASM) is a Database File System that provides cluster file system and volume management capabilities integrated into the Oracle database 10g at no additional cost. ASM lowers your total cost of ownership, increases storage utilization without compromising performance or availability. With ASM, a fraction of the time is needed to manage your database files.

ASM eliminates over provisioning and maximizes storage resource utilization facilitating database consolidation. The ASM self-tuning feature evenly distributes the data files across all available storage. It delivers higher performance than raw, sustained over time, with the ease of use of a file system. ASM's intelligent mirroring technology enables triple data protection, even on non-RAID storage arrays, empowering low cost storage deployment reliably.

ASM reduces Oracle database 10g cost and complexity without compromising performance or availability:

- Simplify and automate storage management
- Increase storage utilization and agility
- Predictably deliver on performance and availability service level agreements


NEC Storage S series products complement ASM and add significant value by implementing the Dynamic pool function leveraging the highly reliable dual parity (RAID6) redundant technology.

The best effective combination of ASM and NEC Storage products can help Database Administrators manage a dynamic database environment by allowing them to grow the database size without having to shutdown the database to adjust the storage allocation. It simplifies the process of setting up a database, adding disks, and removing disks and also configures higher available database system with lower cost.

2. NEC Storage features

2.1 Dual Parity (RAID-6)

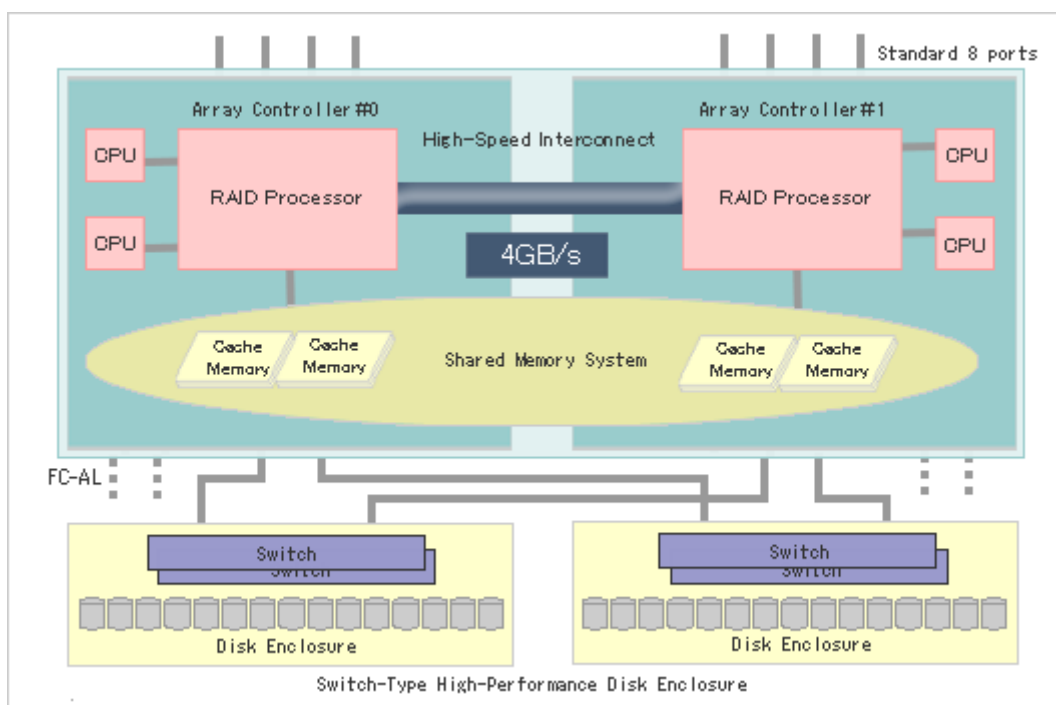
NEC Storage Disk Array is equipped with Dual Parity (RAID-6) technology, which dramatically increases the reliability and availability of the storage. Dual parity mechanism, which is implemented by newly-developed RAID processor, allows continuous operation even if any two hard disk drives (HDDs) fail, thus minimizing recovery time.



Mechanism of RAID-6: Creates two equations based on the data information to generate two independent parity data and record the respective parities (P, Q).

In NEC Storage Disk S4900/S2800/S2400/S1800AT/S1400/S400, the newly developed RAID processor LSI performs this complex RAID operation, and high performance is realized by RAID-6 while maintaining high reliability.

Two HDD Fault Tolerant Technology



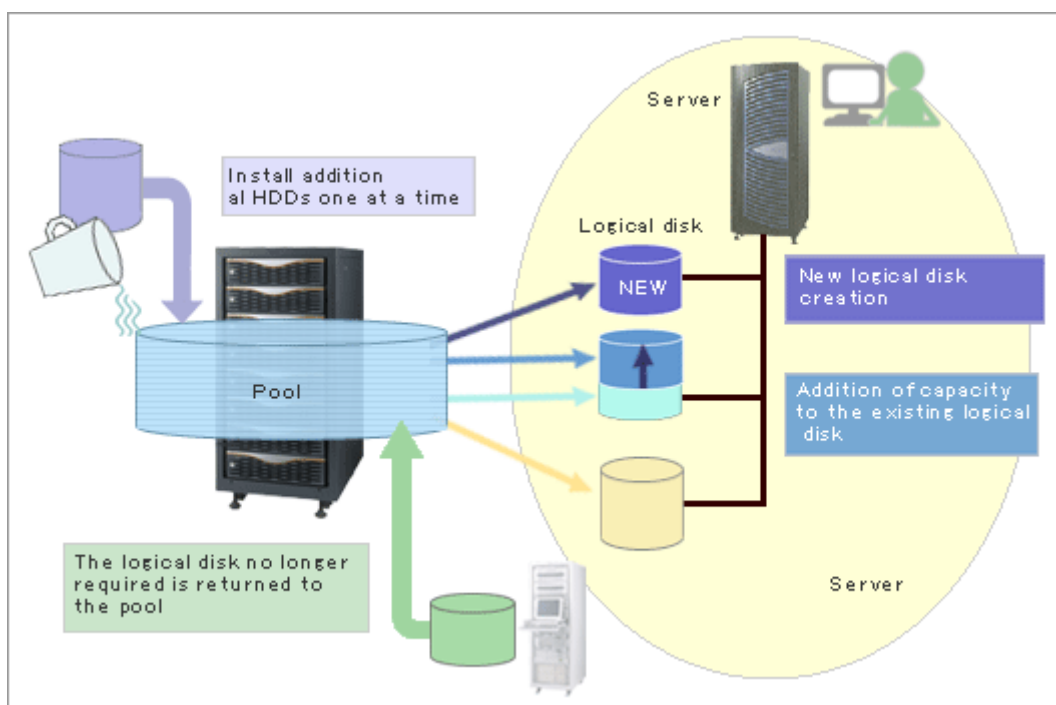
2.2 Dynamic Pool

The dynamic pool function enables flexible and efficient disk capacity management with the virtualization technology.

In addition to the conventional management of disks in units of "RANK," a new dynamic pool function has been developed to enable capacity expansion according to the precise needs, without users being aware of the physical configuration of RAID and RANK (in a RAID-6 configuration).

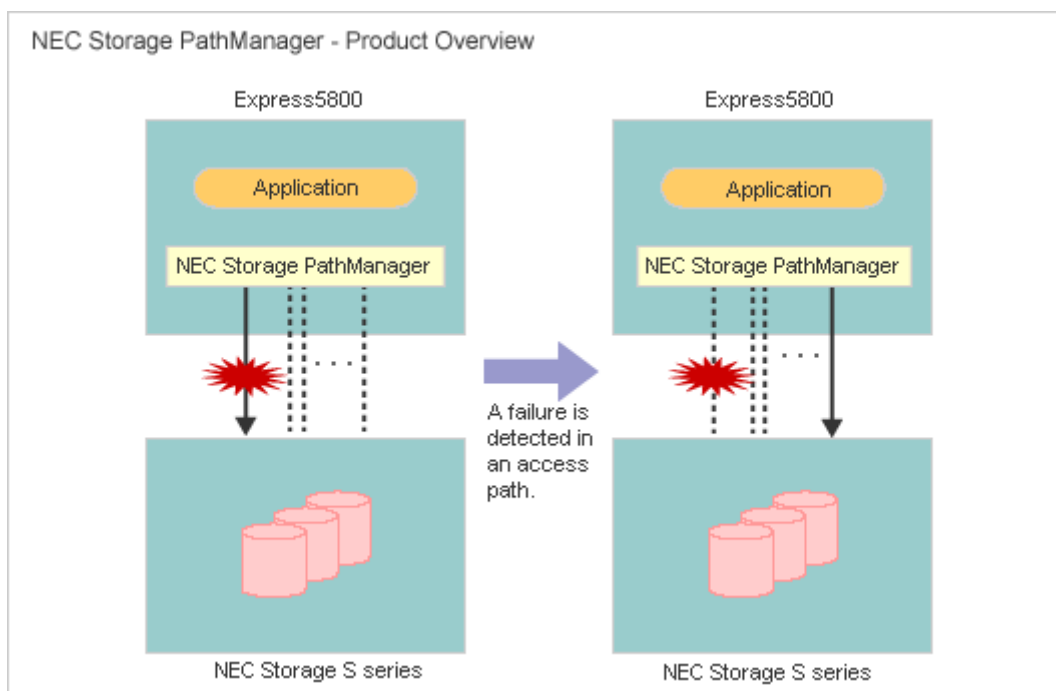
A logical disk with the necessary capacity can be created from the pool. When the pool does not have enough capacity, the capacity can be added by installing additional HDDs one at a time.

When the logical disk becomes unnecessary, it can be separated from the server and returned to the pool. It is also possible to allocate additional logical disk capacity from the pool.



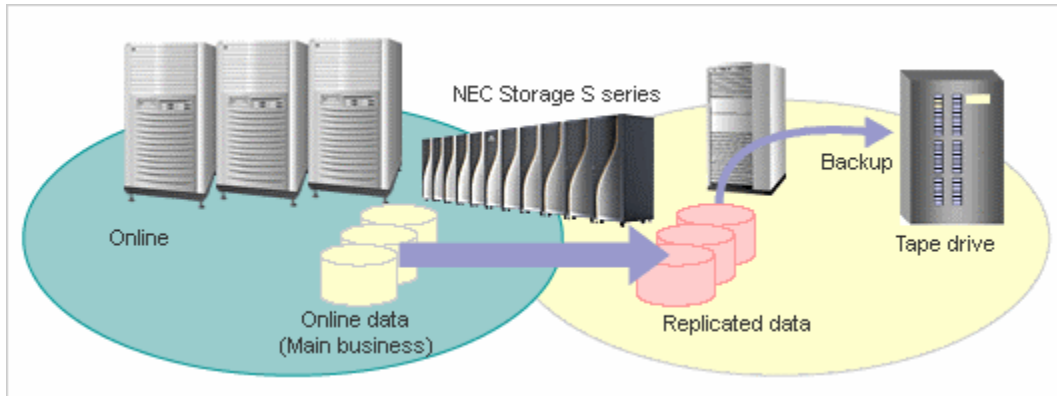
2.3 NEC Path Manager

This software enables automatic switching between access paths when a failure occurs in an access path to NEC Storage Disk Array. After removing the cause of the failure, restoration commands can be used to restore the failed access path. In addition, using multiple access paths at the same time enables balancing of I/O traffic to separate access paths.



2.4 DynamicDataReplication

With DynamicDataReplication, logical volumes in the storage array can be replicated without loading the hosts. These replicated volumes can be split at any point and be used from a host for any other purposes like backups and testing applications.



3. Oracle features

Is it necessary to describe ASM at this level of detail, since we already have so many docs/papers that discuss ASM. You may want to just refer to our BP paper...But its upto you.

3.1 Automatic Storage Management

Automatic Storage Management (ASM) is a vertical integration of both the filesystem and volume manager built specifically for Oracle database files. ASM supports datafiles, logfiles, control files, archive logs, RMAN backup sets, and other Oracle database file types. It spreads database files evenly across all of the disks in a disk group to optimize performance while removing the need for manual I/O tuning. ASM helps DBAs manage a dynamic database environment by allowing them to change the database size without having to shutdown the database to adjust the storage allocation. ASM works with databases running in both exclusive mode and in Real Application Cluster (RAC) environments.

3.1.1 Disk Groups

A *disk group* is a collection of disks managed as a logical unit. ASM spreads each file evenly across all disks in the disk group to balance the I/O. Maximizing load balancing requires that all of the LUNs in a disk group be of similar size and performance and not share common spindles. Most installations will have only a small number of disk groups--for instance, one disk group for a work area and one disk group for a recovery area.

Any single ASM file resides in only a single disk group; however, a disk group may contain files belonging to several databases and a single database may use storage from multiple disk groups. You can specify a disk group as the default disk group for files created in a database by specifying the disk group in file destination initialization parameters.

ASM divides files into 1MB extents and spreads the extents for each file evenly across all of the disks in a disk group. ASM uses pointers to record extent location instead of using a mathematical function to track the placement of each extent. When the disk group configuration changes, ASM can move individual extents of a file rather than having to move all extents to adhere to a formula based on the number of disks.

For files, such as log files, that require low latency, ASM provides fine-grained (128k) striping to allow larger I/Os to be split and processed in parallel by multiple disks. At file creation time, you can decide whether or not to use fine-grained striping. File type specific templates in the disk group determine the default behavior.

3.1.2 The ASM Instance

An ASM instance is a special Oracle instance for coordinating disk group activities. The ASM instance controls the layout of files in a disk group. The ASM instance communicates the file layout information, called an extent map, to database instances when files are created or opened. Once it has the file extent map, a database instance performs I/O directly to disk without involving the ASM instance. During a disk group configuration change, such a disk addition, disk removal, or disk failure, the ASM instance interacts with the database instance to coordinate changes in the file layout.

ASM instances cannot mount a database; they mount disk groups. An ASM instance must be started before a database instance can access files in disk groups. Multiple and separate database instances can share disk groups for their files. On a single node, a single ASM instance typically manages all disk groups. In a Real Application Cluster environment, each node typically has one ASM instance that manages all disk groups for that node in a coordinated manner with the rest of the cluster.

All ASM management commands, such as creating disk groups or adding or dropping disks, must be directed to the ASM instance, and not to the database instance using the ASM Files.

3.1.3 Dynamic Rebalancing

Rebalancing distributes file data evenly across all the disks of the group. ASM automatically rebalances a disk group whenever disks are added or dropped. ASM ensures that a file is evenly spread across all disks in a disk group when the file is allocated, so rebalancing is not required except when the storage configuration changes. With the I/O balanced whenever files are allocated and whenever the storage configuration changes, you never need to search for hot spots in a disk group and manually move data to restore a balanced I/O load. As manual I/O tuning is often a recurring operation in non-ASM environments, ASM's I/O load balancing represents a significant time savings for DBAs.

You can add disks to or drop disks from a disk group that a database is actively using. Once a rebalance completes, added disks carry their share of the I/O load, and you can safely remove dropped disks from the system or reassign them for another purpose. Because the algorithm ASM uses to spread data is not the rigid striping of traditional RAID, ASM does not need to re-stripe all of the data, but instead only needs to move an amount of data proportional to the amount of storage added or removed to evenly redistribute the files and maintain a balanced I/O load across the disks in a disk group.

You can adjust the speed at which data movement occurs during a rebalance. Increasing the speed will make the rebalance complete faster, while reducing the speed will lower the impact on the I/O subsystem.

One example of the power of dynamic rebalancing is the ability to migrate a disk group from an old storage system to a new system while the database is online. When you issue a command to add the disks in the new storage system and drop the disks from the old storage system, ASM automatically moves the data onto the new storage system. When the rebalance is complete, you can safely disconnect the old storage system. Your application does not incur any downtime.

3.1.4 ASM Mirroring

ASM has three disk group types that support different types of mirroring: external, normal, and high redundancy. External redundancy disk groups do not provide mirroring. Use an external redundancy disk group if you use hardware mirroring or if you can tolerate data loss as the result of a disk failure. Normal redundancy disk groups support two-way mirroring or no mirroring on a file by file basis. You can decide at file creation time whether or not to mirror a file in a normal redundancy disk group. File type specific templates determine the default behavior. High redundancy disk groups provide triple mirroring.

ASM uses a unique mirroring algorithm. ASM does not mirror disks, but rather it mirrors extents. As a result, you do not need to have a hot spare disk; you only need spare capacity in your disk group. When a disk fails, ASM automatically reconstructs the contents of the failed disk on the surviving disks in the disk group by reading the mirrored contents from the surviving disks. Not only are all of your spindles active during normal operation, but the I/O hit from a disk failure is spread across several disks, rather than on the single disk that mirrors the failed drive.

When ASM allocates a primary extent of a file to one disk in a disk group, it allocates a mirror copy of that extent to another disk in the disk group. Primary extents on a given disk can have their respective mirror extents on one of several partner disks in the disk group. Each disk in a disk group has the same ratio of primary and mirror extents.

A *failure group* is a set of disks in a disk group sharing a common resource whose failure you wish to protect against. ASM ensures that a primary extent and its mirror copy never reside in the same failure group. If you define failure groups for your disk group, ASM can tolerate the simultaneous failure of multiple disks in a single failure group. This would allow you, for instance, to use ASM to mirror across disk arrays of filers to protect against the failure of an entire disk array or filer.

Since it is tightly integrated with the database, ASM is able to take advantage of database logfiles and other application and file type specific information to eliminate the need for dirty region logging to recover from I/O failures.

4. ASM configuration

This section shows how to create ASM disk groups by using external redundancy of NEC Storage under an Oracle Database 10g Standard Edition RAC environment on a Linux platform.

This paper is written with Linux OS in mind, but the paper and its best practices can be applied to any OS using any releases of Oracle Database 10g.

4.1 Special File names for NEC Path Manager

After installing NEC Path Manager, special files, such as /dev/dda1 as shown below, are created for each Logical Device configured in a NEC Storage in order to encapsulate multiple-access paths to each Logical Device.

```
#ls /dev
/dev/dda1
/dev/dda2
/dev/dda3
/dev/ddb1
/dev/ddc1
/dev/ddd1
/dev/ddd1
/dev/dde1
```

4.2 RAW Device Setting for shared RAC devices

Describe the relationship between Path Manager special file names and RAW devices in "/etc/sysconfig/rawdevices" for Oracle Database to run.

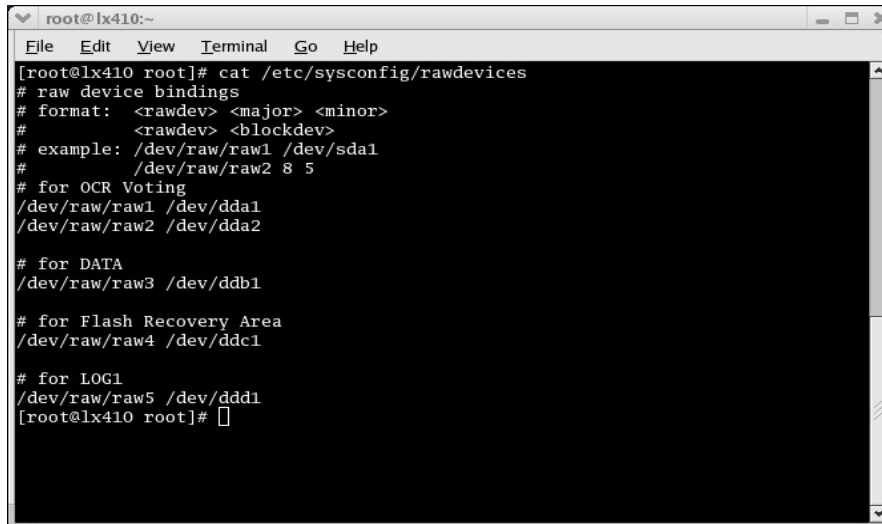
```
Setting sample for /etc/sysconfig/rawdevices

#vi /etc/sysconfig/rawdevices
#for OCR Voting
/dev/raw/raw1 /dev/dda1
/dev/raw/raw2 /dev/dda2
```

```
#for Oracle DATA/SPFILE
/dev/raw/raw3 /dev/ddb1

#for Flash Recovery Area/control1/Archive
/dev/raw/raw4 /dev/ddc1

#for LOG1/control2
/dev/raw/raw5 /dev/ddd1
```



```
root@lx410:~
File Edit View Terminal Go Help
[root@lx410 root]# cat /etc/sysconfig/rawdevices
# raw device bindings
# format: <rawdev> <major> <minor>
#          <rawdev> <blockdev>
# example: /dev/raw/raw1 /dev/sda1
#          /dev/raw/raw2 8 5
# for OCR Voting
/dev/raw/raw1 /dev/dda1
/dev/raw/raw2 /dev/dda2

# for DATA
/dev/raw/raw3 /dev/ddb1

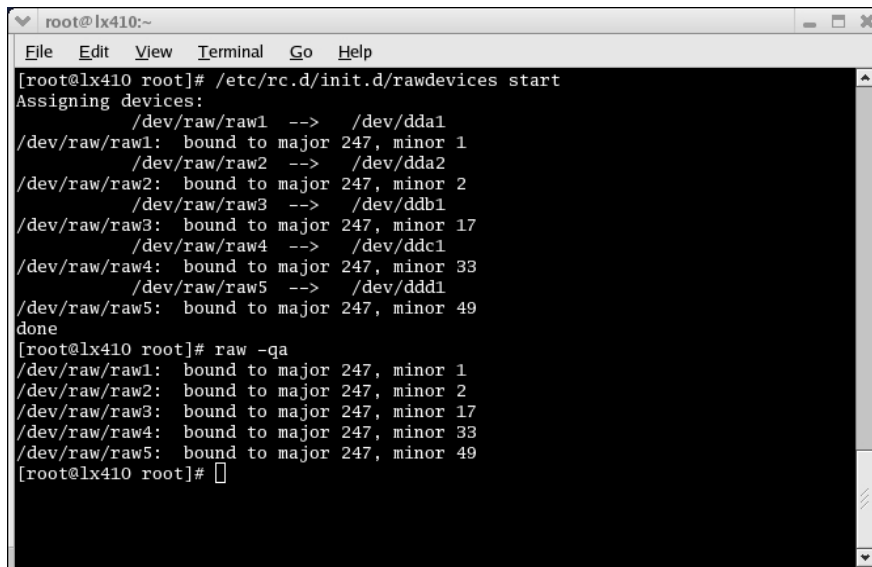
# for Flash Recovery Area
/dev/raw/raw4 /dev/ddc1

# for LOG1
/dev/raw/raw5 /dev/ddd1
[root@lx410 root]#
```

Then, execute "/etc/rc.d/init.d/rawdevices start"

```
Starting rawdevices service

# / etc/rc.d/init.d/rawdevices start
```




```
root@lx410:~
File Edit View Terminal Go Help
[root@lx410 root]# /etc/rc.d/init.d/rawdevices start
Assigning devices:
/dev/raw/raw1 --> /dev/dda1
/dev/raw/raw1: bound to major 247, minor 1
/dev/raw/raw2 --> /dev/dda2
/dev/raw/raw2: bound to major 247, minor 2
/dev/raw/raw3 --> /dev/ddb1
/dev/raw/raw3: bound to major 247, minor 17
/dev/raw/raw4 --> /dev/ddc1
/dev/raw/raw4: bound to major 247, minor 33
/dev/raw/raw5 --> /dev/ddd1
/dev/raw/raw5: bound to major 247, minor 49
done
[root@lx410 root]# raw -qa
/dev/raw/raw1: bound to major 247, minor 1
/dev/raw/raw2: bound to major 247, minor 2
/dev/raw/raw3: bound to major 247, minor 17
/dev/raw/raw4: bound to major 247, minor 33
/dev/raw/raw5: bound to major 247, minor 49
[root@lx410 root]#
```

Setting information for RAW devices can be confirmed by "raw -qa" command.

Next, set owner and group by "chown" command as below for RAW devices to be able to be used by Oracle users.

```
# chown oracle:oinstall /dev/raw/raw[1-5]
```



The above mentioned procedure should be also executed on other servers. Please be sure to use the same RAW device names if RAW devices on each server are bound with the same logical device.

4.3 Creating ASM Disk Groups with External Redundancy

After activating ASM instance, the following command is able to create an ASM disk group with external redundancy implemented by Dual parity (RAID6) feature in NEC Storage Products.

```
SQL> CREATE DISKGROUP DATA EXTERNAL REDUNDANCY DISK '/dev/raw/raw4',
```

Please configure each ASM disk group, depending on RAW device usage sample as follows.

RAW device usage

Raw device name	ASM disk group name	Redundancy	Data
/dev/raw/raw3	+DATA	EXTERNAL	Table space/SPFILE
/dev/raw/raw4	+Flash Recovery Area	EXTERNAL	Online Redo log(primary) Control file(primary)/Archive log file
/dev/raw/raw5	+LOG1	EXTERNAL	Online Redo log(secondary) Control file(secondary)

5. Best Practices for ASM with NEC Storage

This section describes some guidelines to configure ASM disk groups with NEC Storage for running Oracle Database and to backup database online by using the combination of RMAN and ISV's software.

5.1 Disk Groups with Dual parity (RAID6)

To increase high performance and reliability, ASM external redundancy disk groups with NEC's dual parity feature are strongly recommended. Benefits of RAID6 and Dynamic Pool function are:

- Extremely high reliability
- Easy storage configuration
- Efficient storage capacity and provisioning


Below table shows an example of storage configuration using a NEC's entry- level disk array of an S1400.

OS					LD		Pool		
Path(RAW)	Path(RAW)	Path(RAW)	File System	ASM disk group	Name	Resource	No	Physical Disk	RAID Level
/dev/sdb	/dev/dda	/dev/raw/raw1	RAW	Non-ASM	LX_OCR	OCR area	0005h	00h-01h	RAID1
		/dev/raw/raw2				Voting disk			
/dev/sdc	/dev/ddb	/dev/raw/raw3	RAW	+DATA	LX_DATA	DATA/SPFILE	0006h	02h-07h	RAID6
/dev/sdd	/dev/ddc	/dev/raw/raw4	RAW	+Flash Recovery Area	LX_Flash Recovery Area	Redo log1 Control file1/Archive log	0007h	08h-0dh	RAID6
/dev/sde	/dev/ddd	/dev/raw/raw5	RAW	+LOG1	LX_LOG1	Redo log2 Control file2	0007h	08h-0dh	RAID6

This is slightly confusing, because later in the paper you discuss multiple OCR/voting disk in 10gr2, the above table doesn't reflect that. I think you need to state that multiple OCR/votng disk should be used for redundancy. If the above is just an example, then maybe its nit needed and refer to the one later in the paper.

Note:

There is a possibility that OCR and Voting disk are separately placed on plural physical disks when OCR and Voting disk are multiplexed by using OracleClusterware or split into different physical disks on purpose.



It is highly recommended to multiplex the redo log files for availability purposes. Therefore, you can place a copy of the redo log file in the Flash Recovery Area Disk Group along with the archive logs, backup files and temp files. The second multiplexed copy of the redo log file may be placed in a separate ASM Log1 Disk Group. In this scenario, the critical redo log file is protected against unexpected storage failure which may result in a Disk Group not being available for a short while. The NEC Storage S series is a highly reliable and available solution that further protects the Oracle database leveraging its unique Dual Parity feature.

5.2 Use Path Manager with ASM for Improved HA and Performance

NEC Storage S series products even from entry-level equip multiple Fibre Channel ports. NEC Path Manager virtualizes multiple LUNs, which can be identified per Fibre channel port but are actually unique

LUN, as single pseudo device. In order to enhance high availability from the viewpoint of I/O access paths and/or improve I/O performance, the use of NEC Path Manager is strongly recommended.

5.3 Use More Than One LUN per ASM Disk Group for Performance

RAID6 striping of NEC Storage realizes an even distribution of I/O to avoid an I/O bottleneck in the storage layer. ASM can evenly distribute database I/O across all disks in an ASM disk group to avoid an I/O bottleneck in the host layer. These two technologies are complementary and make a perfect fit to accomplish a database system that is automatically tuned for performance. Therefore it is best practice to assign more than one LUN to an ASM disk group.

Using multiple LUNs in the DATA disk group considerably reduces the chances of an I/O bottleneck in the host side.

5.4 Dynamic storage capacity expansion

This section describes Dynamic storage capacity expansion without Oracle database instance interruption under a RAC 10g environment on a Linux platform by using the combination of ASM and NEC Storage pooling function.

Procedure on other platforms may vary.

5.4.1 Adding new Disks into ASM Disk Groups

The following procedures show how to add a newly created LUN derived from Storage pool area into an ASM disk group.

This White Paper describes OS rebooting as a universal way to make a newly created disk dynamically identified from Linux OS. Certainly, there are some means to make it for example:

```
# echo "scsi add-single-device a b c d" >/proc/scsi/scsil
```

However, it does not work on some kernel version. Please make sure that dynamic identification procedures are available or not in your actual Linux environment when you would like to use them.

Stop the Oracle database instance on one node in the RAC environment by a srvctl command to be taken over.

```
[oracle@lx410 oracle]$ srvctl stop instance -d SERAC -i SERAC1 -o immediate
```

Then, reboot the node and confirm to be able to identify the additional disk normally after the rebooting.

```
[root@lx410 root]# shutdown -r now
After the reboot completion
[root@lx410 root]# dmesg | less
```

Modify path setting configuration file (/etc/dualpathrc) of NEC Storage Manager to reflect the additional disk information and activate Storage Manager daemon for path patrol again.

```
[root@lx410 root]# mv /etc/dualpathrc /etc/dualpathrc.bak      rename and store for backup
[root@lx410 root]# mkdd                                       execute mkdd command temporary
Couldn't open /etc/dualpathrc
I try auto setting...
[root@lx410 root]# /etc/rc.d/init.d/dd_daemon restart         activate daemon for path patrol again
```

```
[root@lx410 root]# mkdd
```

Make a partition on the additional disk by using fdisk command.

```
[root@lx410 root]# fdisk /dev/dde          use Path Manager device file name
```

Modify the corresponding table (/etc/sysconfig/rawdevices) of partitions and RAW devices and restart rawdevices service to bind the additional disk. Then, change ownership of the added RAW device to Oracle user.

```
[root@lx410 root]# vi /etc/sysconfig/rawdevices
ex.
/dev/raw/raw6 /dev/dde1          assign the additional disk as /dev/raw/raw6
[root@lx410 root]# /etc/rc.d/init.d/rawdevices restart
[root@lx410 root]# raw -qa      confirm normal setting
[root@lx410 root]# chown oracle:oinstall /dev/raw/raw6
```

Execute the above procedures except making a partition on another node in the RAC environment.

ASM Disk Group Expansion

Log in an operational ASM instance.

```
[oracle@lx410 oracle]$ ORACLE_SID=+ASM1 sqlplus

SQL*Plus: Release 10.2.0.1.0 - Production on Fri Apr 7 16:52:03 2006

Copyright (c) 1982, 2005, Oracle. All Rights Reserved.
Enter user-name: / as sysdba
Connected to:
Oracle Database 10g Enterprise Edition Release 10.2.0.1.0 - Production
With the Partitioning, Real Application Clusters and Data Mining options
SQL>
```

Execute SELECT command to check the total storage capacity (TOTAL_MB).

```
SQL> SELECT * FROM v$asm_diskgroup WHERE NAME='DATA';
GROUP_NUMBER NAME SECTOR_SIZE BLICK_SIZE ALLOCATION_UNIT STATE TYPE TOTAL_MB FREE_MB
-----
2 DATA 512 4096 1048576 MOUNTED EXTERN 269308 265244

SQL>
```

Execute ALTER command specifying the additional RAW device to be configured in an ASM disk group.

```
SQL> alter diskgroup DATA add disk '/dev/raw/raw6';
Diskgroup altered
SQL>
```

Execute SELECT command to check the expanded total storage capacity (TOTAL_MB). You can see the increased storage capacity of 33.6GB in this case.

```
SQL> SELECT * FROM v$asm_diskgroup WHERE NAME='DATA';
GROUP_NUMBER NAME SECTOR_SIZE BLOCK_SIZE ALLOCATION_UNIT_SIZE STATE TYPE TOTAL_MB FREE_MB
-----
2 DATA 512 4096 1048576 MOUNTED EXTERN 303351 299285

SQL>
```

6. Conclusion

Although there is some redundancy between server-based Oracle technology and storage-based NEC technology, the best combination of each strong point results in enabling our common customers to establish a higher reliability/availability, higher performance and scalability, easier manageability database system with lower cost. This joint paper is based on mutual strategic alliance partnership between Oracle and NEC. Both parties will continue to provide the best solutions implemented by both parties' state-of-the-art technologies to joint customers in the future.

The Oracle Database 10g Automatic Storage Management and NEC storage technologies are complementary solutions. The best of breed solutions from both companies offer high reliability and availability, high performance and scalability, easier management and lower cost of ownership. These best practices help our joint customers to implement their solutions faster and with less management overhead. Both Oracle and NEC are committed in providing state of the art technologies and reference architectures to further meet our joint customers' business requirements.

7. Appendix

7.1 Resizing Disks within an ASM Disk Group

The following procedures show how to resize an existing LUN in an ASM disk group by the increased capacity derived from Storage pool area. Please note that these procedures are a little bit complex and your operational mistake might cause a serious problem such as data lost. If you would like to use this dynamic LUN resizing, you have to pay the closest attention.

First, to check current storage capacity, execute SELECT command.

```
[oracle@lx410 oracle]$ ORACLE_SID=+ASM1 sqlplus

SQL*Plus: Release 10.2.0.1.0 - Production on Fri Apr 7 16:52:03 2006

Copyright (c) 1982, 2005, Oracle. All Rights Reserved..
Enter user-name: / as sysdba
Connected to:
Oracle Database 10g Enterprise Edition Release 10.2.0.1.0 - Production
With the Partitioning, Real Application Clusters and Data Mining options
SQL> SELECT * FROM v$asm_diskgroup WHERE NAME='DATA';
GROUP_NUMBER NAME SECTOR_SIZE BLOCK_SIZE ALLOCATION_UNIT_SIZE STATE TYPE TOTAL_MB FREE_MB
-----
3 DATA 512 4096 1048576 MOUNTED EXTERN 269308 265244

SQL>
```

Stop the Oracle database instance which accesses the resized RAW character device.

```
[oracle@lx410 oracle]$ srvctl stop instance -d SERAC -i SERAC1 -o immediate;
```

Stop ASM instance which accesses the resized RAW character device.

```
[oracle@lx410 oracle]$ srvctl stop asm -n lx410 -o immediate;
```

If there is no available space in storage pool area, HDDs can be added to storage pool area dynamically.

Expand an existing LUN size by using the configuration management software tool (NEC Storage Manager).

Unbind the targeted RAW character device.

```
[root@lx410 root]# raw /dev/raw/raw4 0 0
```

Re-read area table of the resized LUN by using fdisk utility.

```
[root@lx410 root]# fdisk /dev/ddb                                Use SPS device file name
The number of cylinders for this disk is set to 4422.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
 1) software that runs at boot time (e.g., old versions of LILO)
 2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)
Command (m for help):w                                          Input "w" key
```

Expand the RAW character device by using fdisk utility, that is, once delete the existing partition and then recreate a new partition. Due to a raw character device, the contents of the deleted partition still remain and the resized partition is able to include the former contents.

Bind again the resized RAW device as a block device.

```
[root@lx410 root]# raw /dev/raw/raw4 /dev/ddb1
```

Restart the stopped ASM instance. And then, confirm that there is no error log in \$ORACLE_BASE/admin/+ASM/bdump/alert+ASM1.log and +DATA disk group are mounted normally.

```
[oracle@lx410 oracle]$ srvctl start asm -n lx410 -o open;
```

Restart the suspended Oracle database instance. And then, confirm that there is no error log in \$ORACLE_BASE/admin/SERAC/bdump/alertSERAC1.log and the Oracle database instance is activated normally.

```
[oracle@lx410 oracle]$ srvctl start instance -d SERAC -l SERAC1 -o open;
```

Log in an operational ASM instance.

```
[oracle@lx410 oracle]$ ORACLE_SID=+ASM1 sqlplus
SQL*Plus: Release 10.2.0.1.0 - Production on Fri Apr 7 16:52:03 2006
```


Copyright (c) 1982, 2005, Oracle. All Rights Reserved..

```
Enter user-name: / as sysdba
Connected to:
Oracle Database 10g Enterprise Edition Release 10.2.0.1.0 - Production
With the Partitioning, Real Application Clusters and Data Mining options
SQL>
```

Execute ALTER command to resize the ASM disk group in the table space.

```
SQL> ALTER DISKGROUP DATA RESIZE DISK DATA_0000 SIZE 300G;
Diskgroup altered
SQL>
```

Execute SELECT command to confirm that resized total capacity (TOTAL_MB).

```
SQL> SELECT * FROM v$asm_diskgroup WHERE NAME='DATA';
GROUP_NUMBER NAME SECTOR_SIZE BLOCK_SIZE ALLOCATION_UNIT STATE TYPE TOTAL_MB FREE_MB
-----
2 DATA 512 4096 1048576 MOUNTED EXTERN 312320 308256

SQL>
```

7.2 Backup and Recovery Using DynamicDataReplication (DDR)


DynamicDataReplication creates replication of business volumes defined separately for each logical disk. The main business volume and its replicated volume are known as Master Volume (MV) and Replicated Volume (RV), respectively. RVs are separated at any time and can be used for multiple purposes such as backup, evaluation, etc. DynamicDataReplication provides fundamental functions to support 24/7 operation and the extension of online business hours.

Combination of ASM and DDR enables to perform hot backups and volume copies without any impacts to servers' performance, thereby answering increasing needs for higher operability and availability.

In a service where 24hour-365day non-stop operation is required, it must be necessary for a customer to back up the database without stopping database or putting load on servers. With NEC Storage, it is possible to answer such needs using combination of ASM and DDR.

7.2.1 System Configuration

(1) Hardware configuration



Replicate data files and log files (archive logs) (RV1/RV2) and mount them using the backup server.

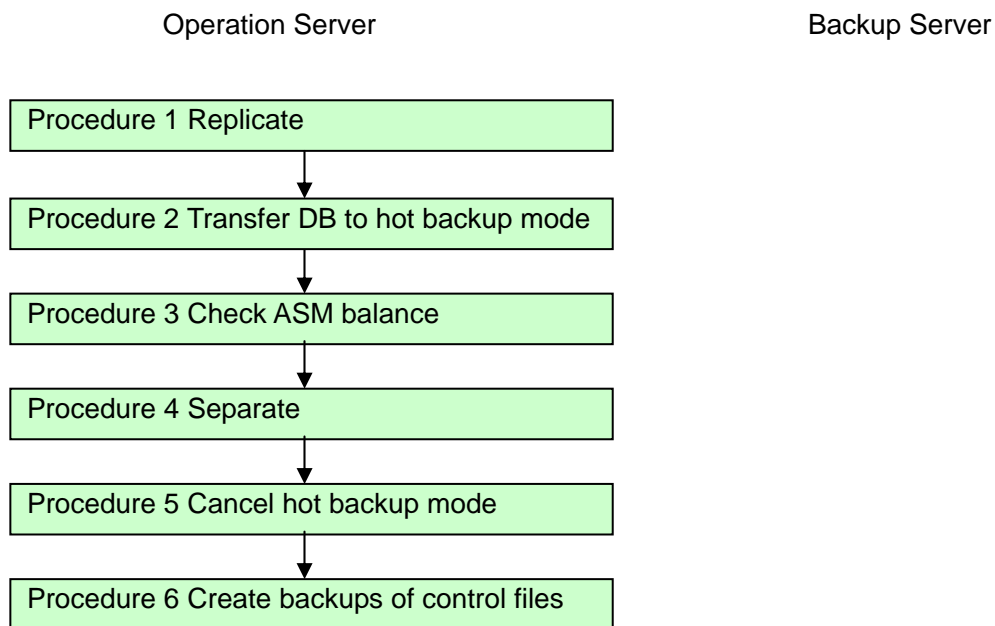
Create backup files of control files and server parameter files in the local drive.

(2) Disk configuration

Access Control	OS			ASM	iStorage LD		iStorage Pool
	Path (SPS)	Path (RAW)	FileSystem	Diskgroup Name	Resource	Capacity	RAID Level
Operation Server (NodeA/NodeB)	/dev/dda1	/dev/raw/raw 11	Raw	-	Voting disk	1000MB	RAID 1
	/dev/dbb1	/dev/raw/raw 21	Raw	-	OCR disk	1000MB	
	/dev/ddc1	/dev/raw/raw 31	Raw	+REDO1	Control file (master)	2000MB	
	/dev/ddd1	/dev/raw/raw 41	Raw		RedoLog file (master)	4000MB	
	/dev/dde1	/dev/raw/raw 51	Raw		Archive log file (master)	20GB	
	/dev/ddg1	/dev/raw/raw 71	Raw	-	Voting disk	1000MB	RAID1
	/dev/ddh1	/dev/raw/raw 81	Raw	-	OCR disk	1000MB	
	/dev/ddi1	/dev/raw/raw 91	Raw	+REDO2	Control file	2000MB	
	/dev/ddj1	/dev/raw/raw101	Raw		RedoLog file (replication)	4000MB	
	/dev/ddk1	/dev/raw/raw111	Raw		Archive log file	20GB	
	/dev/ddm1	/dev/raw/raw131	Raw	-	Voting disk	1000MB	RAID6
	/dev/ddn1	/dev/raw/raw141	Raw	+DATA	Server parameter file	4000MB	
	/dev/ddo1	/dev/raw/raw151	Raw		Table space (MV)	120GB	
		/dev/ddp1	-	EXT3	-	Archive copy area (MV)	20GB
Backup Server	/dev/dda1	-	Ext3	-	Archive copy area (RV)	20GB	RAID1
	/dev/dbb1	-	Raw	-	Table space (RV)	120GB	RAID6

7.2.2 Backup Procedure

(1) Hot backup flow



(2) Hot backup procedure (on production server)

Procedure 1. Replicate (operation server)

Execute replication of the MV where the Oracle table space is stored.

- Execute replication of the MV using the "iSMrc_replicate" command from the console. Wait for the completion of replication using the "-wait" option.

```
[root@operation server]# iSMrc_replicate -file Replication Operation File -wait
```

Contents of the replication operation file

```
0:LX_RAC_DATA_MV
```

```
0:LX_RAC_DATA_RV
```

Procedure 2. Transfer to the hot backup mode **(operation server)**

Notify the database of the start of backup and transfer each table space to the hot backup mode.

- Place all table spaces on the MV to the hot backup mode.

```
SQL> alter tablespace Table Space begin backup; wouldnt database begin backup be better
```

Procedure 3. Check the ASM balance **(operation server)**

This entire Procedure 3 step is not needed if you're using atomic. Please verify that you're using atomic splits.

Check if ASM is being rebalanced.

- **Check if the ASM disk group is being rebalanced by issuing the "select" command. If rebalancing is not being performed, "no rows selected" is displayed.**

```
SQL> select * from v$asm_operation
```

If rebalancing is being performed, rebalancing will be stopped by issuing the "alter diskgroup" command.

```
SQL> alter diskgroup Disk Group Name rebalance power 0
```

Procedure 4. Separate **(operation server)**

Separate the MV and RV synchronized.

- Execute separation by issuing the "iSMrc_separate" command from the console. Wait for the completion of separation using the "-wait" option.

```
# iSMrc_separate -file Replication Operation File -wait
```

Please be sure to use the same replication operation file as the one used for "Replicate."

Procedure 5. Cancel the hot backup mode **(operation server)**

After MV separation is completed, cancel the mode of the table spaces that transferred to the hot backup mode.

- Cancel the mode of all the table spaces that transferred to the hot backup mode.

```
SQL> alter tablespace Table Space end backup; wouldnt database end backup be better
```

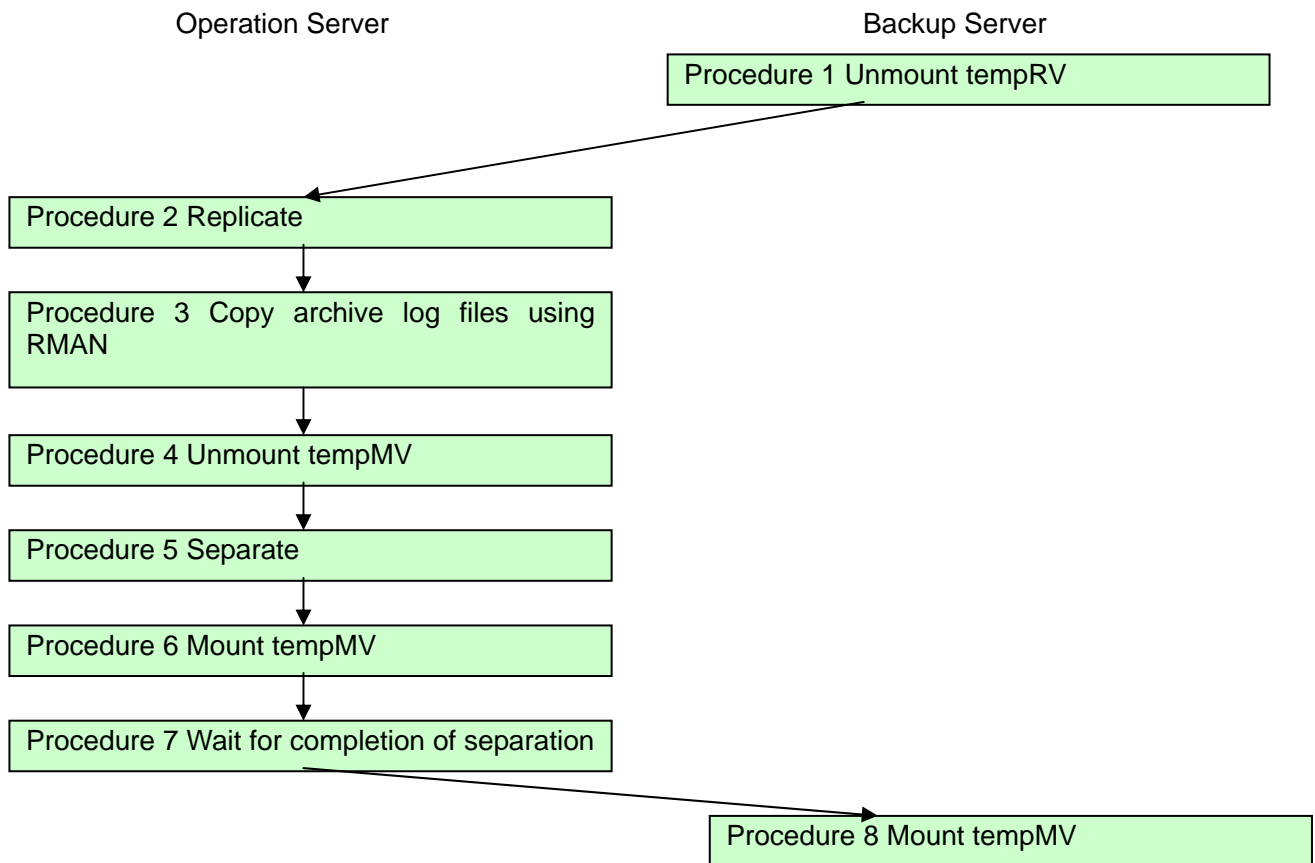
Procedure 6. Create backups of control files **(operation server)**

Create backups of control files and server parameter files of running databases.

```
SQL> alter database backup controlfile to '/u01/app/oracle/ctlSID.bak' reuse;  
SQL> create pfile='/u01/app/oracle/initSID.bak' from spfile=
```

SPFILE Destination Disk Group/spfilename'

(3) Archive Backup Flow



(4) Archive backup procedure (on the backup server)

Procedure 1. Unmount the tempRV (backup server)

Unmount the volume of the RV (tempRV) where copies of archive logs are stored from the system please be specific which server this applies to and reflect the data that remains in the file-system buffer in the disk.

```
[root@backup server]# umount Mount Point
```

Procedure 2. Replicate (operation server)

Execute replication of the volume of the MV (TempMV) where copies of archive logs are stored.

- Execute replication of the MV using the "iSMrc_replicate" command from the console. You can proceed to the next process before the completion of replication using the "-nowait" option.

```
[root@operation server]# iSMrc_replicate -file Replication Operation File -nowait
```


Contents of the replication operation file

```
0:LX_RAC_ARCTMP_MV          0:LX_RAC_ARCTMP_RV
```

Procedure 3. Copy archive log files using RMAN (operation server)

Copy the archive logs in the ASM disk group to the tempMV using RMAN.

- Copy the image copy of the archive logs to the tempMV.

```
RMAN> RUN {  
    BACKUP DEVICE TYPE DISK  
    FORMAT 'tempMV Mount Point!%U'  
    ARCHIVELOG ALL DELETE INPUT};
```

When this command is executed, the current online REDO log file is archived and groups are switched forcibly, and you can proceed with the processing without executing an SQL command.

Procedure 4. Unmount the tempMV (operation server)

Unmount the tempMV from the system and reflect the data that remains in the file-system buffer in the disk.

```
[root@operation server]# umount Mount Point
```

Procedure 5. Separate (operation server)

After completion of the unmounting in Procedure 4, update of the tempMV will be completed. Then execute separation of the tempMV and tempRV.

- By using the RV immediate activation separate function, you can perform separation without waiting for synchronization.

```
[root@operation server]# iSMrc_separate -file Replication Operation File -rvuse immediate -nowait
```

Please be sure to use the same replication operation file as the one used in Procedure 2.

Procedure 6. Mount the tempMV (operation server)

Mount the tempMV unmounted in Procedure 4 on the system.

```
[root@operation server]# mount -t Format Type Device File Name Mount Point
```

Procedure 7. Wait for completion of separation **(operation server)**

The "-nowait" option was used in Procedure 5 by using the RV immediate activation separate function, therefore, separation may not be completed at this point. Wait until the MV and RV are completely separated.

```
# iSMrc_wait -cond sep -file Replication Operation File
```

Please be sure to use the same replication operation file as the one used in Procedure 5.

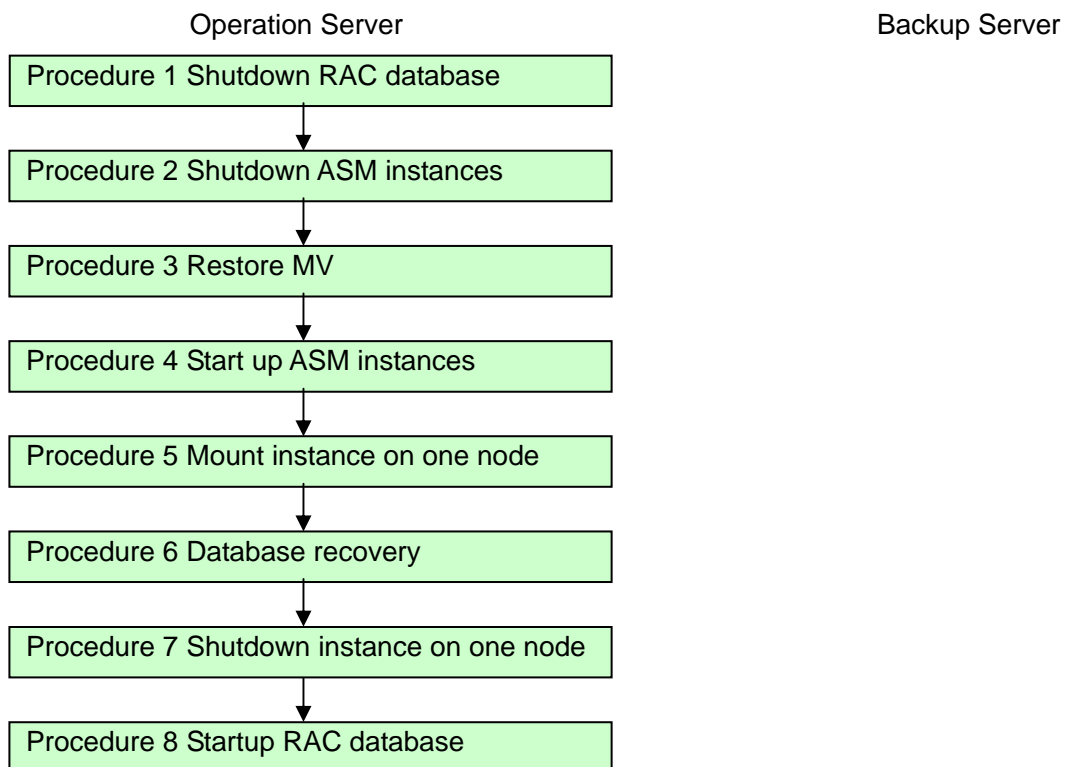
Procedure 8. Mount the tempRV **(backup server)**

After separating the MV and RV in Procedure 5, mount on the system the tempRV that was unmounted in Procedure 1.

```
[root@backup server]# mount -t Format Type Device File Name Mount Point
```

7.2.3 Recovery Procedure

(1) Recovery flow



(2) Recovery procedure

Procedure 1. Shut down the RAC database **(operation server)**

After recovery from a disk failure, etc., stop the RAC database. When the server...please be specific which server this applies to is shutdown, startup the server and stop the RAC database. When Oracle is already stopped, shutdown operation is unnecessary.

- Stop the RAC database using the "srvctl" command.

```
[oracle@operation server]$ srvctl stop database -d Database Name -o immediate
```

Procedure 2. Shut down the ASM instances **(operation server)**

Shut down the ASM instances on both nodes to perform DDR restoration.

- Stop the ASM instances on both nodes using the "srvctl" command.

```
[oracle@operation server]$ srvctl stop asm -n Node Name -i ASM Instance Name -o immediate
```

Please be sure to shut down the instances on both nodes.

Procedure 3. Restore the MV **(operation server)**

Restore the data on the RV to the MV using the data replication function.

- Perform restoration to the MV using the "iSMrc_restore" command from the console.

By using the "-nowait" option and "-mode protect" option, you can proceed to the next step without waiting for synchronization of the MV and RV.

```
[root@operation server]# iSMrc_restore -file Replication Operation File -mode protect -nowait
```

Please be sure to use the same replication operation file as the one used for backup.

Procedure 4. Start up the ASM instances (**operation server**)

Start up the ASM instances on both nodes stopped in Procedure 2.

- Open the ASM instances on both nodes using the "srvctl" command.

```
[oracle@operation server]$ srvctl start asm -n Node Name -i ASM Instance Name -o open
```

Procedure 5. Mount the instance on one node (**operation server**)

Of the instances stopped in Procedure 1, start up the instance used for recovery.

- Mount the database instance using the "srvctl" command.

```
[oracle@operation server]$ srvctl start instance -d Database Name -i Instance Name -o mount
```

Procedure 6. Database recovery (**operation server**)

Perform complete recovery of the database.

```
[oracle@operation server]$ sqlplus / as sysdba
SQL> SET AUTORECOVERY ON;
SQL> SET LOGSOURCE tempMV Mount Point;
SQL> RECOVER DATABASE;
      Execution of Recovery
SQL> ALTER DATABASE OPEN;
```

Procedure 7. Shut down the instance on one node (operation server) why is this required ?

Shut down the instance started up for database recovery.

- Shut down the database instance using the "srvctl" command.

```
[oracle@operation server]$ srvctl stop instance -d Database Name -i Instance Name -o immediate
```

Procedure 8. Start up the RAC database **(operation server)**

Start up the RAC database in the open mode.

- Open the RAC database using the "srvctl" command.

```
[oracle@operation server]$ srvctl start database -d Database Name -o open
```