

Oracle Database 10gTM

Automatic Storage Management with
NEC Storage S Series Products

An Oracle and NEC White Paper

August 2005

Copyright 2005 NEC and Oracle Corporation. All rights reserved.

The content of this document may change without prior notice. This document provides information on the present conditions. This document does not guarantee compatibility with specific purposes such as various conditions or setting examples that are explicitly or inexplicitly indicated in this document, nor does it guarantee compatibility with any other purpose.

Trademarks and registered trademarks

Oracle, Oracle Database 10g, and any other Oracle product name are trademarks or registered trademarks of Oracle Corporation. Linux is a trademark or registered trademark of Linus Torvalds in the United States of America as well as other countries.

Product names mentioned in this document usually are trademarks or registered trademarks of their respective manufacturers/distributors.

This document is a translated version of "Using Oracle Database 10gTM Automatic Storage Management with NEC Storage S Series Products" published by Oracle Corporation/NEC.

Rev. 1.00 August 2005

1. Introduction	4
2. NEC Storage technologies.....	4
2.1 Dual parity (RAID-6)	4
2.2 Dynamic Pool.....	5
2.3 NEC Storage PathManager (SPS)	6
3. Oracle technologies	7
3.1 Automatic storage management	7
4. ASM configuration procedure	10
4.1 Association of SPS special file names.....	10
4.2 Creating settings for RAW devices for a shared disk.....	10
4.3 Creation of an ASM disk group using an external RAID device	13
5. Best practice based on the combination of NEC Storage and ASM	13
5.1 Combination of dual parity (RAID-6) and a disk group	13
5.2 Using NEC Storage PathManager that realizes high reliability and performance	14
5.3 Performance improvement though allocation of multiple LUNs to a disk group.....	15
5.4 Expansion of a dynamic storage capacity	15
6. Conclusion.....	17
7. Appendix	18
7.1 Resizing of an existing logical disk within a disk group.....	18

1. Introduction

This joint White Paper provides information on best practice realized by the combination of the Automatic Storage Management (ASM) of Oracle Database 10g and the NEC Storage S series.

ASM is a database file system included in Oracle Database 10g and provides a cluster file system and the volume management function without requiring any additional costs. ASM can reduce TCO and improve storage availability without lowering its performance or availability. ASM can minimize the time necessary for database file management.

ASM reduces the allocation of excess resources and maximizes storage resource availability to promote database integration. The automatic tuning function of ASM evenly distributes data files to all available storage units. ASM guarantees higher performance than a conventional Raw Device and allows easy system file handling. ASM's intelligent mirroring technology can enable storage array devices, even those without RAID, to protect data using triplex mirroring and procure inexpensive storage.

ASM reduces the cost and complexity of Oracle Database 10g construction with no compromise in performance and availability. In other words, ASM can:

- Simplify and automate storage management,
- Improve storage availability and agility, and
- Provide expected performance and availability based on a Service Level Agreement.


The NEC Storage S series can supplement ASM and provide a valuable addition by using the Dynamic Pool function with the redundant technology of super high-reliability dual parity (RAID-6).

The most effective combination of ASM and NEC Storage can allow the expansion of database size without the administrator having to shut down the database for the purpose of adding storage, and this will then allow the administrator to dynamically manage the database. This combination can simplify the processes of data base setup, disk addition, and disk deletion, and also allows the construction of high-availability databases with lower costs.

2. NEC Storage technologies

2.1 Dual parity (RAID-6)

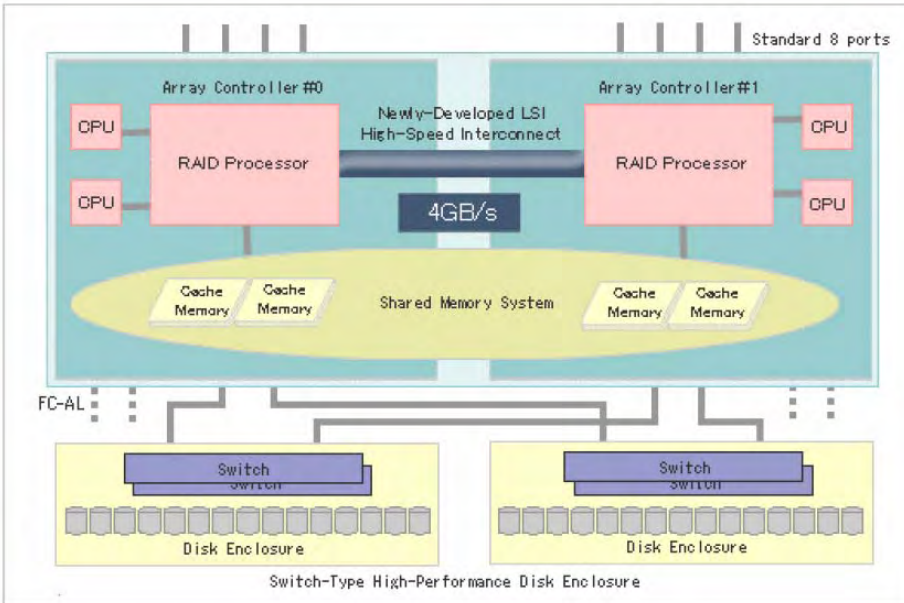
The NEC Storage S series is equipped with dual parity (RAID-6) technology that dramatically improves reliability and availability. The dual parity mechanism, realized with a new RAID processor uniquely developed by NEC, allows operations to continue even if two HDDs fail at the same time. In other words, this mechanism can minimize the time necessary for restoration.



Two HDD Fault Tolerant Technology

RAID-6 mechanism: based on the stored data, this mechanism generates two equations, creates two pieces of independent parity data, and records each of them (P and Q).

NEC Storage S4900/S2800/S2400/S1800AT/S1400/S400 use a newly developed RAID processor to execute this complicated RAID operation on the hardware. As a result, it has become possible to achieve high performance while maintaining high reliability.



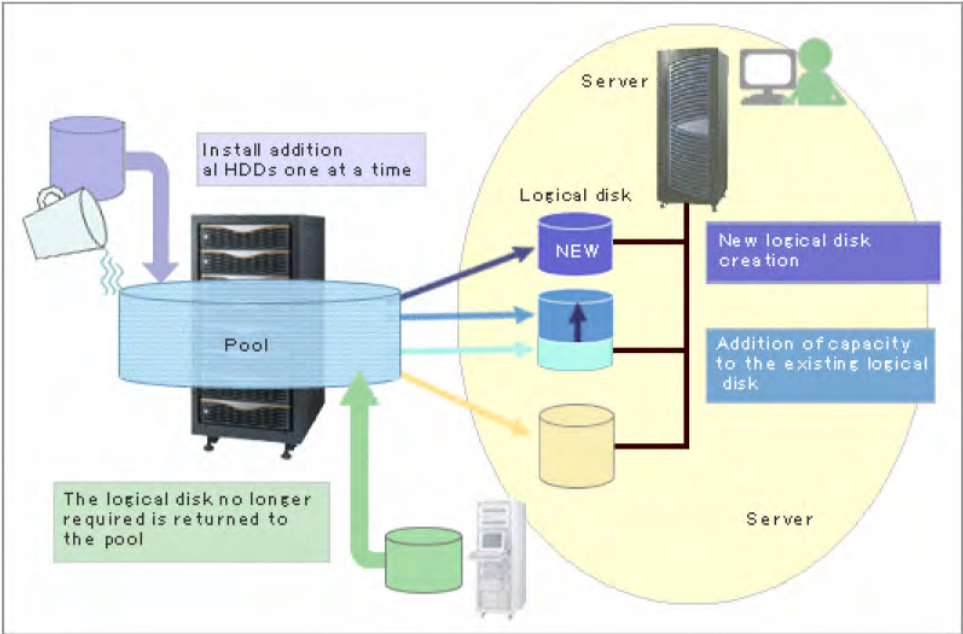
2.2 Dynamic Pool

The Dynamic Pool function uses virtualization technologies to enable flexible and efficient disk management.

In addition to the conventional disk management method, in which disks are managed in units of "RANK," the Dynamic Pool function has been developed to allow disk capacity expansion in accordance with the correct capacity required for a disk, without our customers being conscious of physical configurations of RAID and RANK when creating a RAID-6 configuration.

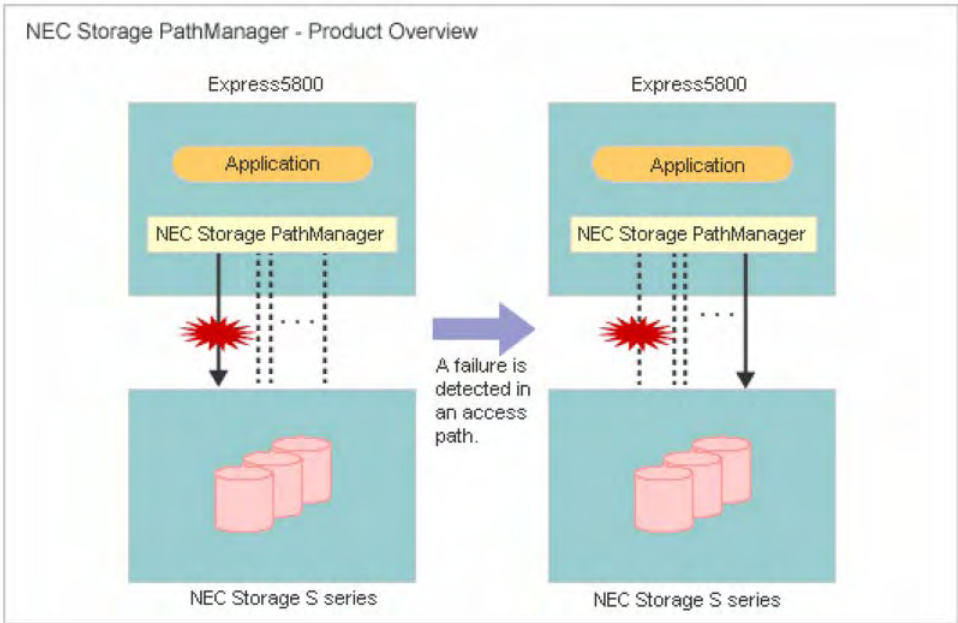
From the "pool," it is possible to create a logical disk having the necessary capacity. If the pool does not have sufficient capacity, then capacity can be added to the pool in units of HDD.

When a logical disk is no longer necessary, it can be separated from the server and returned to the pool. It is also possible to allocate the necessary capacity from the pool to increase the volume of an existing logical disk.



2.3 NEC Storage PathManager (SPS)

This software automatically switches access paths when a failure occurs in the access path to NEC Storage. To restore the access path after the cause of failure is removed, the "restoration" command is used. Furthermore, by using multiplex access paths, it is possible to distribute I/O traffic to other paths.



3. Oracle technologies

3.1 Automatic storage management

Automatic Storage Management (ASM) is realized by the vertical integration of a file system, which is designed specifically for Oracle database files, and a volume manager. ASM supports Oracle database files such as data files, log files, control files, archive logs, and RMAN backup sets. Database files are evenly allocated to all disks within the disk group; therefore, it is possible to optimize performance and eliminate the need for manual I/O tuning. With ASM, DBA can dynamically change database size without shutting down the database in order to adjust a storage area. ASM can be used in Real Application Clusters (RAC) environments and in non-RAC configuration databases.

3.1.1 Disk group

A disk group is a group of disks that is treated as one logical unit. ASM allocates individual files to all disk groups within the disk group to balance I/O loads. In order to maximize load balancing, it is necessary that all LUNs in the disk group have similar size and performance, and do not share the spindle. A small number of disk groups are used in most installations. For example, one disk group for a working area, and the other disk group for a recovery area.

Any single ASM file can exist only in a single disk group. However, one disk group can contain files that belong to multiple databases, and one database can use areas of multiple disk groups. The default disk group of a database file can be specified as a database initialization parameter "file destination."

ASM divides a file into 1 MB extents and evenly allocates them to all disks within the disk group. ASM uses a pointer instead of a formula to track these extents. When the disk group configuration changes, ASM can move the individual file extents without using a formula and thus displacing them.

For files requiring low latency (such as log files), ASM uses a smaller (128 k) strip size such that a request with a large I/O size can be divided and multiple disks can process them in parallel. Whether or not striping should be in a fine grain size can be determined at the time of file creation. The default value is determined by the template that is dependent on the type of a file within the disk group.

3.1.2 ASM instance

An ASM instance is a special Oracle instance for synchronizing disk group operations. The ASM instance controls the file layout within the disk group. When creating or opening a database file, the ASM instance passes layout information called an extent map to a database

instance. Once the file extent map is obtained, the database instance executes I/O directly to the disk without ASM instance intervention. The ASM instance synchronizes database instances and then changes in file layout while the disk group configuration is being changed (while disks are being added or deleted, or when a failure has occurred, for example).

An ASM instance cannot mount a database, but it mounts a disk group. An ASM instance must be initiated before the database instance accesses files within the disk group. Different database instances can share a disk group. In a single-node configuration, a single ASM instance manages all disk groups. In a Real Application Clusters environment, an ASM instance exists at each node, tuned within a cluster, and manages all disk groups at that node.

ASM management commands (disk group creation, or addition or deletion of disks) are all issued to an ASM instance and not to a database instance that uses ASM files.

3.1.3 Dynamic rebalancing

Rebalancing refers to even allocation of file data to all disks within a disk group. ASM automatically rebalances the disk group when disks are added or deleted. ASM operates in such a way that files are evenly allocated to all disks within a disk group, and therefore rebalancing is required only when a storage configuration is changed. Since the I/O balance is adjusted at the time of file allocation and storage configuration changes, it is no longer necessary to search for a hotspot within the disk group and manually move data in order to distribute I/O loads. In a non-ASM environment, manual I/O tuning is often repeated, and I/O loads distribution by ASM means significant time saving for the DBA.

Disks can be added to or deleted from a disk group during database operations. When rebalancing is complete, the added disks receive I/O loads, and deleted disks can be removed from the system or used for different purposes. Since the algorithm used for data allocation by ASM is not the conventional and strict RAID striping, ASM does not have to re-stripe all data. The volume of data to be moved to reallocate data and create an I/O load balance is proportional to the volume of added or deleted disks.

The speed of data movement during rebalancing can be adjusted. Rebalancing ends earlier as the speed is increased, and when the speed is decreased, it is possible to reduce the impact on the I/O subsystem.

One of the advantages of dynamic rebalancing is that disk groups can be migrated from an old storage system to a new system while the database is online. When a command is issued ordering the addition of disks to a new storage system and the removal of disks from an old storage system, ASM automatically moves data to the new storage system. After rebalancing is complete, the old storage system can be separated. There will be no downtime in your applications.

3.1.4 ASM mirroring

ASM supports three types of mirroring for disk groups: external, regular, and high redundancy mirroring. Disk groups with external redundancy do not use mirroring. Mirroring is used when hardware uses mirroring or when data loss due to disk failure can be tolerated. Disk groups with regular redundancy support either bidirectional mirroring or no mirroring for each file. Whether or not files within disk groups with regular redundancy should be mirrored is determined at the time of file creation. A file type-specific template determines the default value. Disk groups with high redundancy support three-way mirroring.

ASM uses a unique mirroring algorithm. It does not mirror disks; it mirrors extents. Therefore, there is no need to have a hot spare disk ready. A spare area is necessary only within the disk group. In the case of disk failure, ASM reads the mirrored content from remaining disks to automatically reconstruct the content of the failed disk on these remaining disks. All the spindles are active during regular operations, and an I/O spike triggered by disk failure is distributed to multiple disks in addition to the disk that had mirrored the failed disk.

When ASM secures the primary extent of a file on a particular disk within a disk group, a mirror copy of that extent is secured in a different disk in that disk group. The corresponding mirror extent of the primary extent on a particular disk exists on a particular disk in the disk group. All disks in that disk group have the same volume of primary extents and mirror extents.

A failed group refers to a group of disks within a single disk group that share the same resources to be used in the case of failure. ASM guarantees that it will never place the primary extents and their mirror copies in the same failure group. When a failure group is defined for a disk group, ASM can tolerate simultaneous failures of multiple disks within a single failure group. For this reason, when arrays of filer groups are mirrored, it is possible to be ready for failure of an entire filer.

ASM is closely integrated with databases. It uses database log files, other applications, or file type-specific information to eliminate the necessity for Dirty Region Logging for recovering from I/O failure.

4. ASM configuration procedure

This section explains how to create an ASM disk group using the RAID function of NEC Storage in the Oracle Database Standard Edition RAC environment on a Linux platform.

This White Paper is written for Linux, but the descriptions of the best practices can be applied to other types of OS, other release versions of Oracle Database 10g.

4.1 Association of SPS special file names

When NEC Storage PathManager (SPS) is installed, the following special files are created for each logical device configured within NEC Storage in order to hide multiple paths to the logical devices.

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

4.2 Creating settings for RAW devices for a shared disk

Relationship between RAW devices, which are the RAC shared disks, and SPS special file names are described in the "/etc/sysconfig/rawdevices" file. In the example below, "/dev/raw/raw1" to "/dev/raw/raw5" are bound to the RAW devices.

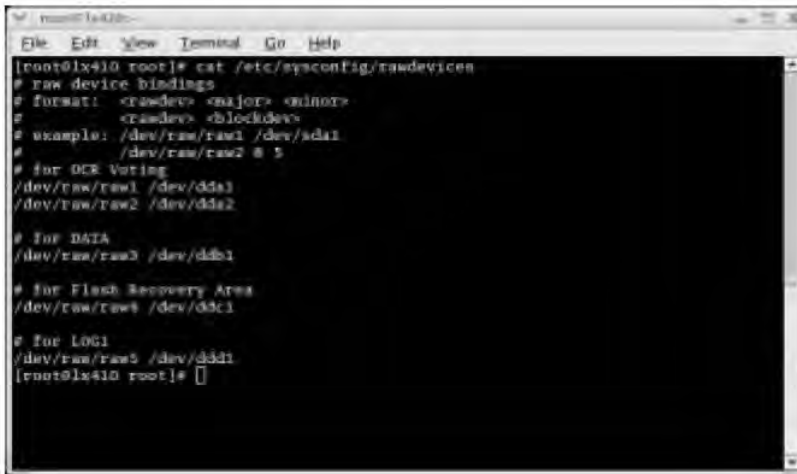
```
Setting sample for /etc/sysconfig/rawdevices
```

```
#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

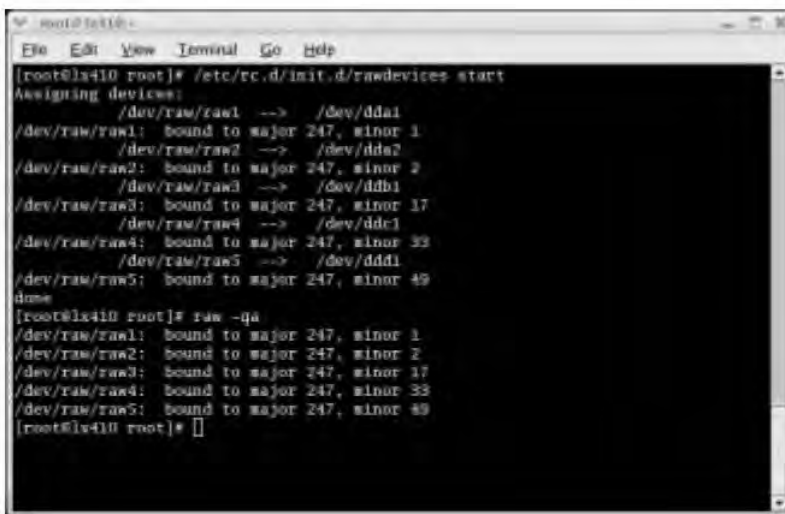
#for LOG1/control2
/dev/raw/raw5 /dev/ddd1
```



```
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 # 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]#
```

Next, the "rawdevices" service is started.

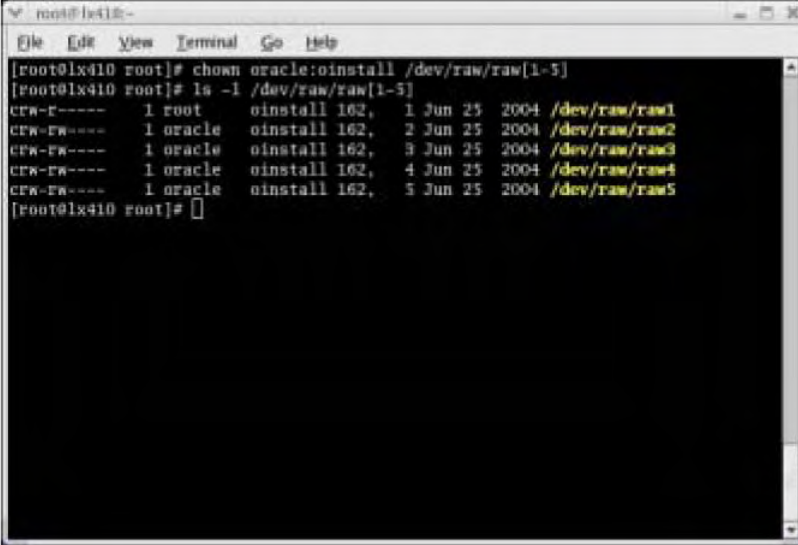
```
Starting rawdevices service
# / etc/rc.d/init.d/rawdevices start
```



```
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]#
```

Settings can be checked with the "raw-qa" command. In the screen shown above, each disk has been found as specified in setting files. Settings for the owner and group are then necessary to enable an Oracle user to use the created RAW devices. Settings can be made with the "chown" command.

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



```
root@lx410:~# chown oracle:oinstall /dev/raw/raw[1-5]
root@lx410:~# ls -l /dev/raw/raw[1-5]
crw-r----- 1 root    oinstall 162, 1 Jun 25 2004 /dev/raw/raw1
crw-rw---- 1 oracle  oinstall 162, 2 Jun 25 2004 /dev/raw/raw2
crw-rw---- 1 oracle  oinstall 162, 3 Jun 25 2004 /dev/raw/raw3
crw-rw---- 1 oracle  oinstall 162, 4 Jun 25 2004 /dev/raw/raw4
crw-rw---- 1 oracle  oinstall 162, 5 Jun 25 2004 /dev/raw/raw5
root@lx410:~#
```

Create the same settings for other servers configuring RAC. The important point here is using the same name to bind disks having the same physical properties for all the servers.

4.3 Creation of an ASM disk group using an external RAID device

After starting an ASM instance, use the command shown below to use an external RAID device to create an ASM disk group. Since the NEC Storage S series supports dual parity (RAID-6), it is possible to construct database storage with extremely high reliability.

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

Refer to the table below for disk applications. Execute the above command for the number of disk groups to be used.

RAW device applications

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 practice based on the combination of NEC Storage and ASM

This section provides an ASM disk group configuration guideline for online backups through functional links between Oracle database operations, ISV backup software, and Oracle RMAN (Recovery Manager).


5.1 Combination of dual parity (RAID-6) and a disk group

It is strongly recommended that an ASM disk group using dual parity (RAID-6) of NEC Storage be created in order to realize both high performance and availability at the same time. Also, the combination of RAID-6 and the Dynamic Pool function has the following advantages:

- Very high reliability
- Simple storage configuration
- Efficient use and provisioning of storage capacity

The table below shows examples of storage configurations using S1400, which is the entry model of the NEC Storage S series.

S					D		ool		
Path(RAW)	ath(RAW)	ath(RAW)	file system	SM disk group	ame	esource	o	hysical disk	AID 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



It is also strongly recommended that Redo logs be made redundant for availability improvement. More specifically, replications of Redo logs should be placed in the Flash Recovery Area disk group together with archive logs, backup files, and temporary files. It is also possible to place replication of Redo logs in the ASM Log1 disk group. In this scenario, critical Redo logs can be protected from unexpected short-term disk group failure. The NEC Storage S series is a solution with high reliability and availability that protects Oracle databases by using its uniquely developed dual parity function.

5.2 Using NEC Storage PathManager that realizes high reliability and performance

All the NEC Storage S series products, from the entry level models to higher level models, are equipped with redundant fiber channel ports.

The entity of NEC Storage PathManager is a single LUN, but it virtualizes the LUN, which is recognized as multiple LUNs via multiple paths, as a single pseudo-device. From the

perspective of reliability and performance improvement of I/O access paths, the use of NEC Storage PathManager is strongly recommended.

5.3 Performance improvement through allocation of multiple LUNs to a disk group

The striping function of RAID-6 of NEC Storage evenly allocates I/O within the storage layer to prevent I/O becoming a bottleneck. ASM evenly allocates I/O among all disks within a disk group to prevent I/O becoming a bottleneck within the host layer. These two technologies supplement each other and realize a database system in which performance is automatically tuned. Therefore, the best practice is to allocate multiple LUNs to an ASM disk group.

By allocating multiple LUNs to a DATA disk group, it is possible to dramatically reduce the possibility of I/O creating a bottleneck in the host.

5.4 Expansion of a dynamic storage capacity

This section describes the procedures for dynamic storage capacity expansion without stopping database instances by using the combination of ASM and the pool function of NEC Storage in the Oracle Database 10g SE RAC environment on the Linux platform.

Note that the procedures are different for other OS platforms.

5.4.1 Dynamic addition of a logical disk to a disk group

This section describes the procedures for creating a new LUN using a storage pool area and adding it to an ASM disk group.

In this White Paper, OS reboot is used as a universal method for making the Linux OS recognize a newly created logical disk. There is of course a command like the following to make the OS recognize the disk:

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

However, this may not function properly depending on the kernel version. Therefore, if you plan to use a dynamic recognition method other than OS reboot, make sure to check beforehand whether that method is effective in the Linux environment actually in use.

Next, use the "srvctl" command to stop the Oracle database instances on one node, which is part of the RAC configuration, to make a standby system node take over the operations.

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

Reboot the node. Check that the newly added logical disk has been recognized.

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

In order to incorporate the addition of the new disk, recreate the status setting file (/etc/dualpathrc) to be used by NEC Storage PathManager. Then restart the path patrol daemon.

```
[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
```

Use the "fdisk" command to create a partition in the added logical disk.

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

Define the correspondence between the created partition and a RAW device within the "etc/sysconfig/rawdevices" file. Then restart the "rawdevices" service to bind the disk. Also, change the owner to an Oracle user since the added RAW device is used with Oracle software.

```
[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
```

Except for the creation of a partition, follow the same procedures for the standby system server so that the disk will be available for use.

ASM Disk Group Expansion

Log into an ASM instance at the node actually in use.

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

SQL*Plus: Release 10.1.0.3.0 - Production on Fri Apr 22 13:08:51 2005

Copyright (c) 1982, 2004, Oracle. All rights reserved.
Enter user-name: / as sysdba
Connected to an idle instance.
SQL>
```

Execute the following SELECT statement and check the current total volume shown as 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 the following ALTER statement to specify a RAW device of a disk to be added to an ASM disk group "+DATA" in a table space.

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

Execute the following SELECT statement to check the post-expansion total volume shown as TOTAL_MB. In this example, a 33.6 GB disk is newly added.

6. Conclusion

There are overlapping technologies between server-based Oracle technologies and storage-based NEC technologies, but by combining their advantages in the best way, it becomes possible to provide customers shared by NEC and Oracle with higher reliability, availability, performance, scalability, and better database manageability at lower cost. This joint White Paper is based on the strategic alliance partnership between Oracle and NEC. Oracle and NEC will continue to provide their customers with the best solution using their most advanced technologies.

ASM and NEC storage technologies are supplemental solutions. The best-of-breed solutions of the two companies can provide high reliability, high availability, high performance, high scalability, database manageability, and reduction of ownership costs. Such a best practice will help the customers obtain the solutions earlier with a lower management overhead. Oracle and NEC are striving to provide the most advanced technologies and reference architectures in order to further fulfill the customer's business needs.

7. Appendix

7.1 Resizing of an existing logical disk within a disk group

This section describes the procedures for resizing the existing LUN within an ASM disk group by taking the volume necessary for addition out from the storage pool area. Since these procedures are somewhat complicated, operational mistakes may cause serious problems such as data loss. Therefore, follow the procedures with the utmost caution.

First, execute the SELECT command to check the current capacity of the disk to be expanded.

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

SQL*Plus: Release 10.1.0.3.0 - Production on Wed Apr 27 11:27:32 2005

Copyright (c) 1982, 2004, Oracle. All rights reserved.
Enter user-name: / as sysdba
Connected to an idle instance.
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>
```

In order to expand the RAW character device, stop the Oracle instance that is accessing the device.

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

In the same manner, stop the ASM instance that is accessing the device.

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

If there is no available capacity in the storage pool area, it is possible to expand the pool capacity by dynamically adding an HDD.

Using NEC Storage Manager (iSM) client, expand the size of the existing LUN.

Unbind the target RAW character device.

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

Use the fdisk utility to reload the space table of the expanded LUN.

```
[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
```

Use the fdisk utility to expand the RAW character device. Delete the current partition and create a new one. Due to the nature of the RAW character device, it is possible to create a new partition with an expanded capacity without deleting the content even though the old partition is deleted.

Re-bind the RAW character device with an expanded volume to the block device.

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

Restart the ASM instance. Check that error messages are not displayed in "\$ORACLE_BASE/admin/+ASM/bdump/alert+ASM1.log" and that the +DATA disk group is properly mounted.

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

After the ASM instance starts up normally, start the Oracle instance. Then, check that error messages are not displayed in "\$ORACLE_BASE/admin/SERAC/bdump/alertSERAC1.log" and that the Oracle instance started normally.

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

Log into an ASM instance on the server actually in use.

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

SQL*Plus: Release 10.1.0.3.0 - Production on Wed Apr 27 11:27:32 2005

Copyright (c) 1982, 2004, Oracle. All rights reserved.
Enter user-name: / as sysdba
Connected to an idle instance.
SQL>
```

Expand the capacity of the disk group "+DATA" in the table space. Execute the ALTER statement to resize the disk group.

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

Execute the SELECT statement to make sure that the TOTAL_MB value has increased.

```
SQL> SELECT * FROM v$asm_diskgroup WHERE NAME='DATA';
GROUP_NUMBER
-----
NAME
-----
SECTOR_SIZE BLOCK_SIZE ALLOCATION_UNIT STATE
-----
TYPE          TOTAL_MB  FREE_MB
-----
```

```
DATA
   512   4096   1048576 MOUNTED
EXTERN   312320  308256
SQL>
```