

Best Practices White Paper

**Using Oracle Database 10g
Automatic Storage Management with FUJITSU Storage**

Sep 6, 2005
Fujitsu Limited

Contents

1. Introduction	1
2. Fujitsu's ETERNUS Storage Systems	2
2-1. High Availability.....	2
2-2. Performance Tuning.....	2
2-3. Point-in-time copy feature.....	2
3. Fujitsu Storage Management software	2
3-1. Fujitsu Storage Cruiser.....	2
3-2. Fujitsu AdvancedCopy Manager.....	3
4. Oracle Technologies	3
4-1. Automatic Storage Management.....	3
4-2. ASM Architecture.....	3
4-3. I/O Distribution.....	3
4-4. Dynamic Reconfiguration of a disk group.....	4
4-5. Rebalancing.....	4
4-6. ASM Mirroring.....	4
4-7. ASM Recovery.....	4
5. Best Practices	4
5-1. High Reliability and High Availability.....	4
5-2. Database Acceleration.....	4
5-3. System Relation Management and Failure Monitoring.....	5
5-4. ASM Disk Group Configuration Guideline.....	5
6. Conclusion	6
7. Appendix	7
7-1. Cluster System by the Combination of Oracle RAC and PRIMECLUSTER.....	7
7-2. Database Space Management by Fujitsu Storage Manager.....	8
7-3. Backup/Recovery Procedure Methods and Usage Example.....	9
7-4. Additional Technical References.....	24

1. Introduction

This document describes the most efficient database environment that provides the combined benefits and advantages of the new Oracle Database 10g feature, Automatic Storage Management (ASM), and Fujitsu ETERNUS storage systems.

Automatic Storage Management is a new feature in Oracle Database 10g that provides data management features within the database that simplifies the management of database systems and storage environments. ASM is the combination of a clustered file system and logical volume manager that is tightly integrated and embedded within the Oracle database 10g. ASM was specifically designed to operate with Oracle database files. ASM is a standard Oracle Database 10g feature that is included at no additional cost with the database system. ASM is designed to optimize the operation of Oracle databases and simplify storage management by virtualizing storage hardware, automating tasks and optimizing performance thereby reducing operational overhead, administrator intervention – and human error – and overall management costs.

ASM is Oracle's collection of intelligence and best practices learned over many years. The ASM disk group virtualization interface allows DBAs to easily and dynamically add or remove disks for a particular disk group without interrupting the database operation. ASM will automatically begin to use a new disk as soon as it is added to the disk group, and will automatically rebalance data between old and new disks to continually optimize I/O performance. ASM also eliminates disk fragmentation and hotspots thereby reducing management overhead.

ASM provides up to 3-way mirroring for higher availability and is an alternative to mirroring provided by the hardware RAID in storage arrays. The Oracle DBA simply defines a failure groups within a disk group and ASM will automatically mirror the data across two or more failure groups, which can help to avoid downtime in the event of component failures, such as a disk controller or even entire disk array. This powerful redundancy feature increases availability without any additional management overhead. Since failure groups are automatically managed by ASM fail over or fail back activities are fully managed without the intervention of an operator by ASM when the failed component is replaced.

Fujitsu's ETERNUS is the storage systems that meet the requirements in mission-critical-environment. ETERNUS6000 is highly scalable, large capacity disk array, and ETERNUS3000 is mid range disk array. In any model, all major components are redundantly equipped and are hot-swappable to ensure business continuity with high reliability and high availability. ETERNUS storage systems provide an online volume-replication feature that enables non-disruptive database backup operation.

This white paper describes the most efficient method to configure a high-performance, high-availability Oracle database environment by the combination of ASM and ETERNUS.

2. Fujitsu's ETERNUS Storage Systems

2-1. High Availability

Fujitsu ETERNUS storage system provides high availability that guarantees secure operation in mission critical environments.

All major components are redundantly equipped. In the event of a component failure, ETERNUS automatically reallocates the rest of active components for non-disruptive operation. Plus, all major components are hot-swappable.

Disk drives are configured in RAID groups. In the event of a disk-drive failure, the data is quickly rebuilt in a hot-spare disk to maintain the data redundancy. Plus, ETERNUS storage system constantly monitors each disk drive for pre-failure detection. If one of the disk drives is likely to fail, ETERNUS proactively copies the data to the hot-spare disk. This enables disk-drive exchange while maintaining the data redundancy.

The interfaces between servers and ETERNUS storage systems also deliver a high available mechanism. This mechanism is provided by ETERNUS Multipath Driver or GR Multipath Driver that manages multiple Fibre Channel paths. Even if one path fails, operations of a database and other applications can continue by automatically rerouting data through the rest of active paths. Such a highly available mechanism guarantees continuous operations even in the event of a server-side HBA failure or switch failure.

2-2. Performance Tuning

Fujitsu ETERNUS storage system provides features to exchange drives or to change RAID configurations without disrupting accesses to the volumes. With these features, it is possible to optimize access performance without stopping operations.

RAID Migration provides two features to reorganize disk drives. One is a feature to include the currently-used RAID drives into the destination RAID group. This feature is useful when reconfiguring the same-specification drives into a different RAID-level group. The other feature is useful when moving the data from the source RAID drives to a different type of target RAID drives, for example, from a 10,000rpm-drive RAID group to a 15,000rpm-drive RAID group. Both features can be used without quitting accesses to the RAID group, and the switching from the source to destination is automatically performed by ETERNUS.

2-3. Point-in-time copy feature

Fujitsu ETERNUS storage system provides a point-in-time copy feature that creates a copy of the source volume data at a certain point in time within the storage system without stopping the operations of a database or other applications.

System requirement for 24 x 365 continuous operation is growing. System administrators must make backups without disrupting business operations, such as database operation, and without putting heavy loads on the servers. ETERNUS provides two types of replication features: Equivalent Copy (EC) and One Point Copy (OPC). With EC, the source and target volumes are previously synchronized and split when the replicated volume is required. With OPC, a copy volume can be instantly created on demand. Both of these features can be used according to business requirements.

3. Fujitsu Storage Management software

3-1. Fujitsu Storage Cruiser

Fujitsu Storage Cruiser offers the advanced storage resource management features based on the SAN (Storage Area Network) management features. With versatile management features for configuration, troubleshooting, performance and correlation, Fujitsu Storage Cruiser provides stable operation for any type of storage system, including SAN, DAS (Direct Attached Storage) and NAS (Network Attached Storage).

In SAN environments, the management of only physical connection topologies is not enough. Without the management of logical connection topologies, it is impossible to exactly understand the impact of a device/component failure. Fujitsu Storage Cruiser not only graphically displays logical connections, but also quickly reflects reconfigurations or component failures on the display. This enables a quick action in the event of a system failure.

3-2. Fujitsu AdvancedCopy Manager

Fujitsu offers Fujitsu AdvancedCopy Manager, an application to control the Equivalent Copy and One Point Copy features of ETERNUS. With Fujitsu AdvancedCopy Manager, it is possible to manage replicated spaces, and to non-disruptively create a consistent database copy synchronizing with the original database operation.


4. Oracle Technologies

4-1. Automatic Storage Management

ASM enables a disk group to be designated as storage for Oracle database files, control files and backup files. A disk group consists of multiple disk drives and is managed as a unit by ASM. Any file stored in a disk group is automatically striped over all the disk drives in the group. ASM provides clustering capabilities and is available not only for single SMP machine but across multiple nodes of Oracle Real Application Clusters. In order to generate maximum performance, ASM automatically and dynamically distributes I/O loads among all the disk drives, even if data usage pattern is rapidly changing. It prevents fragmentation of space so there is no need to relocate data for shortage of space. By selecting the desired reliability and performance level of operation, ASM sets the management characteristics for each file, and does not require operator's interaction on every file.

4-2. ASM Architecture

ASM feature is performed by 'ASM instance' that is different from 'DB instance'. ASM instance manages data mapping information in a disk group by using metadata describing the layout of ASM files.



DB instance performs actual access to data files by using the data mapping information from the ASM instance.

4-3. I/O Distribution

A file (an ASM file) created on a disk group is divided evenly and stored among all the disk drives in the disk group. This prevents I/O concentration on a single disk drive. There are two stripe widths: COARSE and FINE. The default for control file and redo log file is FINE and for others it is COARSE.

4-4. Dynamic Reconfiguration of a disk group

When a disk drive is added to or removed from a disk group or if a disk drive fails, ASM automatically redistributes data among all the disk drives. This operation can be performed without interrupting database operations.

4-5. Rebalancing

When a disk group is reconfigured, the data is reallocated so that the load on each disk in the disk group becomes even. This operation is called 'rebalancing'. Rebalancing has some performance impacts on online transactions because it generates internal accesses for moving data. Rebalancing does not relocate data based on I/O statistics nor is it started as a result of statistics. It is completely driven by the size of the disks in the disk group and is automatically started when the storage configuration changes.

4-6. ASM Mirroring

ASM has three disk group types that support different types of mirroring, first one is normal redundancy, second one is high redundancy and last one is external redundancy. Normal redundancy disk groups support two-way mirroring. High redundancy disk groups support three-way mirroring. External redundancy disk groups do not provide mirroring. When RAID storage is used, we recommend external redundancy disk group.

4-7. ASM Recovery

If the ASM instance shuts down or ends abnormally, the DB instance depending on the ASM instance also ends abnormally. This, however, will not crash the operating system and therefore can be recovered and restarted quickly and efficiently.

With an RAC configuration database, an ASM instance is distributed to each node. This enables that even if an ASM instance on a certain node ends abnormally, the ASM instances on the other nodes automatically recovers and maintains the operation.

5. Best Practices

5-1. High Reliability and High Availability

ASM in Oracle Database10g provides a mirroring feature that further enhances Oracle database availability. However, for entire availability enhancement, ETERNUS's high-availability features including its RAID mechanism are the optimum solution. ETERNUS can perform maintenance operations including failure detection, retry by switching and replacement of failed parts concurrently with database operations.

5-2. Database Acceleration

Database performance can be enhanced by the combination of the ASM striping feature and the RAID mechanism. The ASM striping feature allocates data across all the disks in the disk group to enhance the entire throughput. Moreover, with the disk group volume configured as RAID0+1 or RAID5, the performance is enhanced by striping within the ETERNUS.

Database operations sometimes require performance tuning after the operations have started. ASM optimizes file accesses by distributing file extents to all the disks in a disk group. This means manual tuning is not required.

Moreover, ETERNUS provides RAID migration feature that enables performance tuning from another aspect. With RAID migration, it is possible to change disk drives comprising the RAID group, for example, from 10,000rpm drives to 15,000rpm drives and to change RAID levels, for example, from RAID5 to RAID0+1, without stopping operations.

5-3. System Relation Management and Failure Monitoring

If a system has component redundancy, operations will not stop even in the event of a component failure. However the failure could affect the performance. To minimize the impact of such a system component failure, it is required to detect the failed part and quickly understand the extent of the impact. The combination of ASM and Fujitsu Storage Cruiser enables the centralized relation management of database instances, access paths on the SAN and physical drives in ETERNUS.

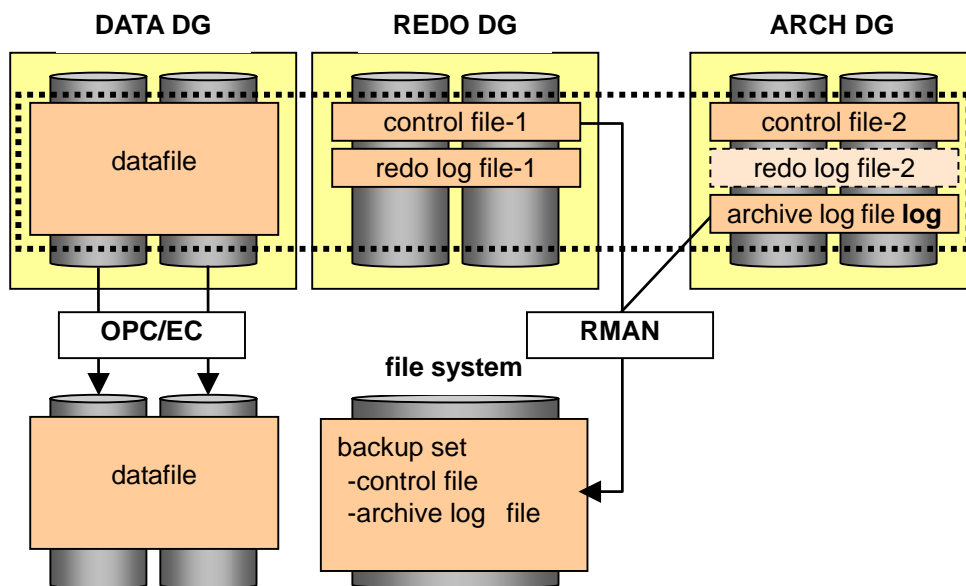
5-4. ASM Disk Group Configuration Guideline

Datafile is huge in size, and does require concurrent backup and rapid recovery, so it is best to use Advanced Copy for backup of datafile. On the other hand, control file and archive log file are backed up with recovery manager. There is no need to back up redo log file. It is best to allocate each file in separate disk group for ease of operation at the time of restoration. Especially, if Advanced Copy is used for datafile backup, then redo log file should never be placed in the same disk group as datafile, because Advanced Copy will overwrite all the data of the disk group at the time of restore.

In order to maximize the benefit of ETERNUS Advanced Copy feature, we recommend to configure ASM disk groups as described in the figure below.

- (1) Create three disk groups. The first one is for the datafile (DATA DG), the second one is for the redo log file (REDO DG) and the last one is for the archive log file (ARCH DG).
- (2) Allocate only the datafiles to DATA DG.
- (3) Allocate all the redo log files and the control file to REDO DG.
- (4) Allocate all the archive log files and a mirror of the control file to ARCH DG.
- (5) Datafile is backed up by ETERNUS Advanced Copy feature (OPC or EC). On the other hand, control file and archive log file are backed up by recovery manager.

Note: If you want to maximize the availability of redo log files, then you can place a mirror of all redo log file logs in the archive log file disk group (ARCH DG).



As for the ETERNUS RAID configuration, RAID0+1 is recommended for online volumes considering performance and reliability, and RAID5 is recommended for backup volumes considering cost-effectiveness. If the database capacity is small or the performance requirement is not too demanding, RAID1 could also be an option for online volumes. For the maximum performance, it is recommended that the disk group device is configured as a 1partition=1LUN=1RAID configuration.

6. Conclusion

To configure a high-performance, high-availability and excellent-operability database environment with the combination of the new Oracle 10g ASM feature and Fujitsu ETERNUS storage system, the following points must be considered.

- (1) Configure disk groups with external redundancy.
- (2) Employ RAID0+1 or RAID5 for volumes to be registered in a disk group.
- (3) Create three disk groups, DATA DG, REDO DG and ARCH DG. Allocate only datafiles on DATA DG, allocate all redo log files and a control file to REDO DG, and allocate all archive log files and a mirror of a control file to ARCH DG.
- (4) For maximum availability of redo log file, it is recommended to place a mirror of redo log file in archive log file disk group.
- (5) RAID Migration feature enables performance tuning by changing the rotational speed of disk drive and/or the RAID levels concurrently with online operations.
- (6) With the relation management feature of Fujitsu Storage Cruiser, it is possible to centrally manage relations between database instances, associated disk groups, the names of the devices comprising disk groups, access paths on the SAN and the physical drives of ETERNUS. This provides a quick understanding of the extent of the impact in the event of a disk failure.


7. Appendix

As described above, this Best Practices explains the basic features available by the combination of Oracle 10g ASM, ETERNUS and Fujitsu Storage Cruiser. For higher-reliability and higher-operability requirements, we recommend considering the following solutions.

7-1. Cluster System by the Combination of Oracle RAC and PRIMECLUSTER

To configure a higher-reliability/availability system by the combination of Oracle10g ASM and ETERNUS, we recommend a cluster system by the combination of Oracle10g RAC and Fujitsu PRIMECLUSTER.

With the cooperation of Oracle10g RAC and PRIMECLUSTER, it is possible to switch not only Oracle but also all the online applications to another server in the event of a server failure.



7-2. Database Space Management by Fujitsu Storage Manager

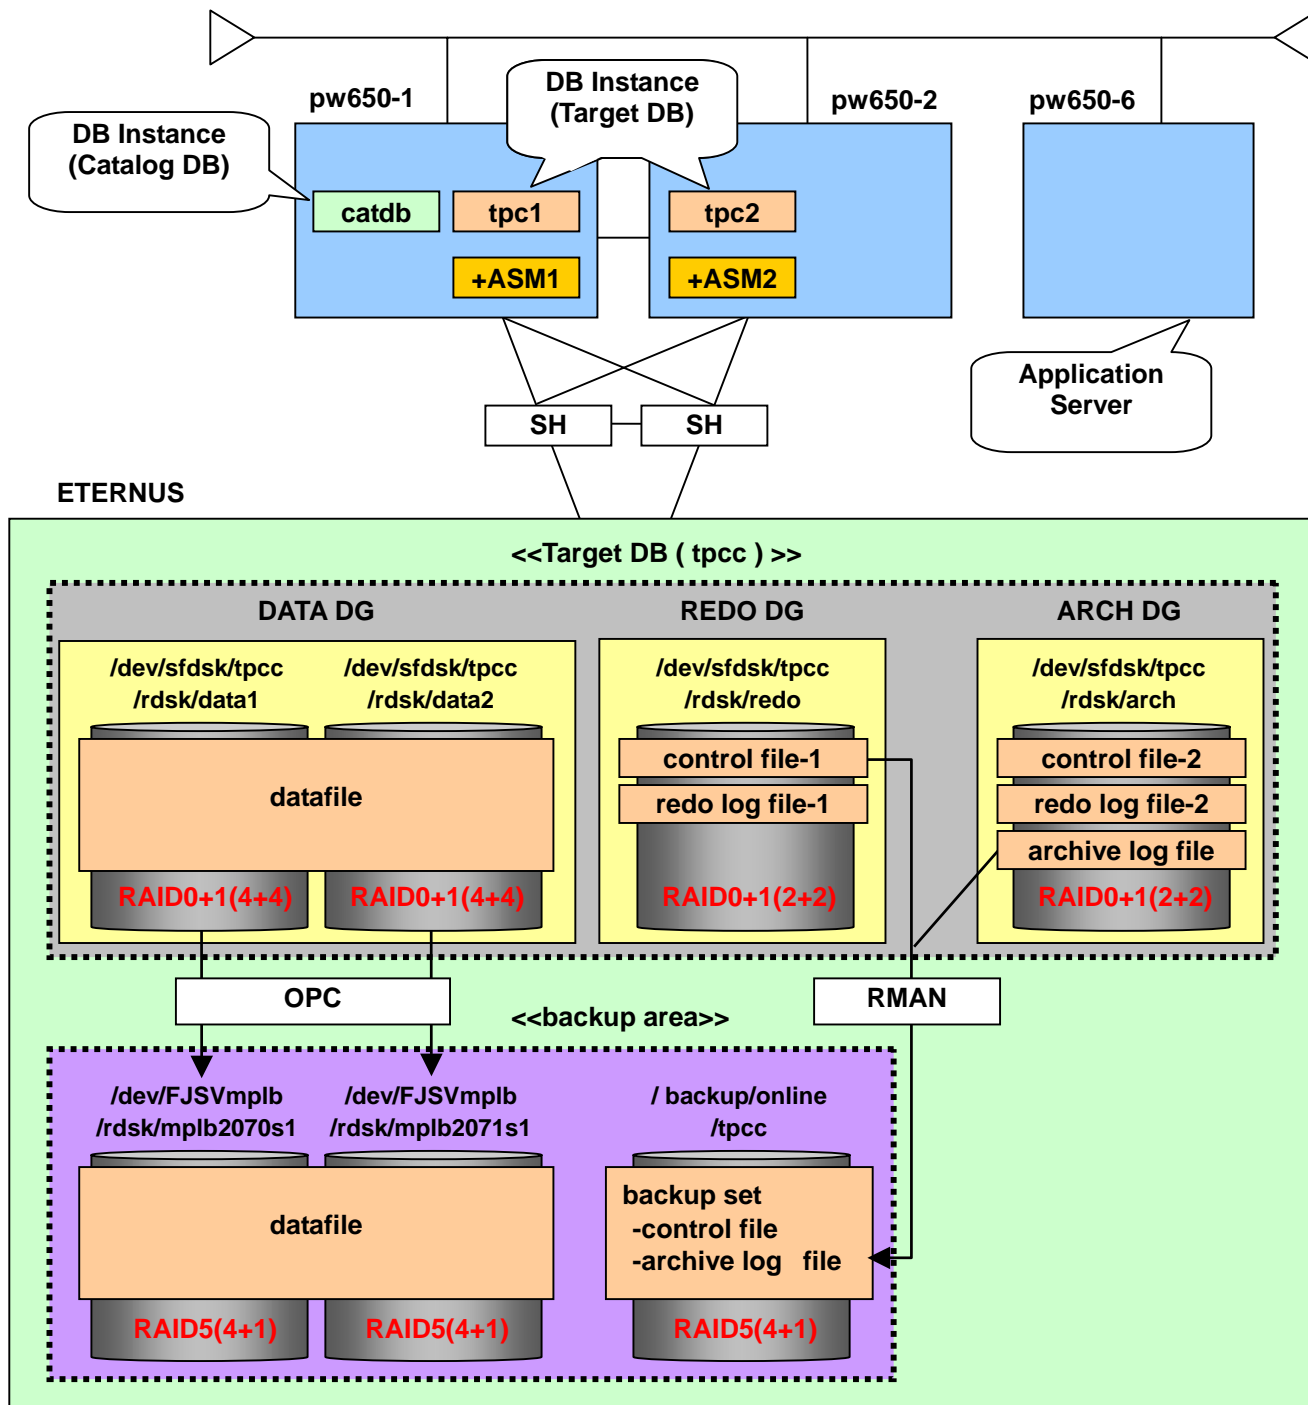
Fujitsu Storage Manager provides storage-space management. Fujitsu Storage Manager enables centralized storage-space management for a wide variety of storage-environment types including SAN, DAS (Direct Attached Storage) and NAS (Network Attached Storage). The features consist of monitoring of space usage, policy-based space management and usage-status analysis. These features support stable system operations and help reduce the total cost of ownership (TCO).

By adding Fujitsu Storage Manager to the combination of Oracle10g ASM and ETERNUS, not only the visible areas of Oracle 10g ASM on ETERNUS but also the space information of the system disk in the internal disk drives can be centrally managed. This delivers the following benefits:

- Prevention of a system down by the capacity monitoring
- Effective capacity planning by the space-usage-ratio analysis and the prediction report
- Understanding of the DB re-designing period against performance degradation due to capacity growth

7-3. Backup/Recovery Procedure Methods and Usage Example

This section provides step-by-step procedures for performing backup and recovery of ASM file system using Fujitsu Advanced Copy feature with ETERNUS storage system.



7-3-1. Backup/Recovery Procedure Methods and Usage Example

1. Archive a current redo log file group.

```
SQL> alter system archive log current;
```

2. Place the database in backup mode.

```
SQL> alter database begin backup;
```

3. Ensure that no ASM rebalance operations are occurring.

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

This should report no rows. If a row is returned; i.e., a rebalance occurring, then stop the rebalance operation using the following:

```
SQL> alter diskgroup DG_DATA rebalance power 0;
```

And re-check v\$asm_operation.

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

4. Back up all disks belonging to a disk group of datafile using Advanced Copy feature.
See Also : <7-3-3 Backup Procedure Usage Example>

5. Take the database out of backup mode.

```
SQL> alter database end backup;
```

6. Back up all archive log files. A current control file is backed up using the auto backup feature of control file.

```
RMAN> backup archivelog all;
```

If a server parameter file exists, that needs to be converted to a text based file, in preparation for cases that a disk group containing the sever parameter file gets corrupted. Fujitsu strongly suggests to set RMAN CONFIGURE CONTROLFILE AUTOBACKUP that automatically backs up the server parameter file.

Whenever a datafile is added to the database, it is necessary to immediately perform the database back up.

7-3-2. Recovery Procedure Methods

When a datafile disk group is broken.

1. Shutdown an instance of database with the ABORT option.

```
SQL> shutdown abort
```

2. Dismount a disk group of datafile. In case of RAC, it is necessary to dismount this disk group from all the nodes of the RAC. In some cases the diskgroup will get dismounted.

```
SQL> alter diskgroup DG_DATA dismount;
```

3. Restore all disks belonging to a disk group of datafile using Advanced Copy feature.
See Also : <7-3-4 Recovery Procedure Usage Example>

4. Mount a disk group of datafile.

```
SQL> alter diskgroup DG_DATA mount;
```

5. Startup an instance of database with the MOUNT option.

```
RMAN> startup mount;
```

6. Perform media recovery with a complete recovery.

```
RMAN> recover database;
```

7. Open the database.

```

RMAN> alter database open;

```

When all disk groups are broken.

1. Shutdown an instance of database with the ABORT option.

```

SQL> shutdown abort

```

2. Dismount a disk group of datafile. In case of RAC, it is necessary to dismount this diskgroup from all the nodes of the RAC.

```

SQL> alter diskgroup DG_DATA dismount;

```

3. Drop the other disk groups. In case of RAC, it is necessary to dismount the disk group from all the nodes of the RAC before dropping the disk group.

```

SQL> drop diskgroup DG_REDO including contents;
SQL> drop diskgroup DG_ARCH including contents;

```

4. Initialize disk header of all disks belong to DG_REDO, and DG_ARCH.

```

$ dd if=/dev/zero of=/dev/sfdsk/tpcc/rdsk/redo bs=1024 count=10
$ dd if=/dev/zero of=/dev/sfdsk/tpcc/rdsk/arch bs=1024 count=10

```

5. Create a disk group of DG_REDO and a disk group of DG_ARCH.

```

SQL> create diskgroup DG_REDO external redundancy disk
  2 '/dev/sfdsk/tpcc/rdsk/redo';
SQL> create diskgroup DG_ARCH external redundancy disk
  2 '/dev/sfdsk/tpcc/rdsk/arch';

```

6. Restore all disks belonging to a disk group of datafile using Advanced Copy feature.
See Also : <7-3-4 Recovery Procedure Usage Example>

7. Mount a disk group of datafile.

```

SQL> alter diskgroup DG_DATA mount;

```

8. Copy a backup control file to work area.

```

$ cp /backup/online/tpcc/c-1844893648-20050409-00 /backup/ct100.bak

```

9. Startup an instance of database without mount.

```

RMAN> startup nomount;

```

10. Restore the control file from a control file at work area.

```

RMAN> restore controlfile from '/backup/ct100.bak';

```

11. Mount the database.

```

RMAN> alter database mount;

```

12. Perform media recovery with an incomplete recovery.

```

RMAN> recover database until sequence 127 thread 1;

```

13. Open the database with the RESETLOGS option.

```

RMAN> alter database open resetlogs;

```

14. Add any temporary files to a temporary tablespace. No need to specify path name, keep it an OMF name.

```

SQL> alter tablespace TEMP add tempfile
  2 '+DG_DATA/tpcc/temp01.dbf' size 100m reuse;

```

15. Delete some invalid entries of archive log files.

```

RMAN> crosscheck archivelog all;
RMAN> delete expired archivelog all;

```

7-3-3. Backup Procedure Usage Example

**Archive a current redo log file group.
Perform this command on database instance of pw650-1.**

```
SQL> alter system archive log current;
```

System altered.

```
SQL>
```

**Place the database in backup mode.
Perform this command on database instance of pw650-1.**

```
SQL> alter database begin backup;
```

Database altered.

```
SQL>
```

**Ensure that no ASM rebalance operations are occurring.
Perform this command on ASM instance of pw650-1.**

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

no rows selected

```
SQL>
```

**Back up all disks belonging to a disk group of datafile using
Advanced Copy feature.
Perform this command on OS command line of pw650-1.**

```
# /opt/FJSVswsts/bin/swstbackup /dev/sfdsk/tpcc/dsk/data1:mplb2052
/dev/sfdsk/tpcc/dsk/data1:mplb2052 swstbackup completed
# /opt/FJSVswsts/bin/swstbackup /dev/sfdsk/tpcc/dsk/data2:mplb2054
/dev/sfdsk/tpcc/dsk/data2:mplb2054 swstbackup completed
#
```

**Take the database out of backup mode.
Perform this command on database instance of pw650-1.**

```
SQL> alter database end backup;
```

Database altered.

```
SQL>
```

**Back up all archive log files. A current control file is backed
up using the auto backup feature of control file.
Perform this command on recovery manager of pw650-1.**

```
RMAN> backup archivelog all;
```

```
Starting backup at 05/08/09
current log archived
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=133 devtype=DISK
channel ORA_DISK_1: starting archive log backupset
channel ORA_DISK_1: specifying archive log(s) in backup set
input archive log thread=1 sequence=123 recid=16 stamp=565811250
input archive log thread=1 sequence=124 recid=17 stamp=565811456
```

```

input archive log thread=1 sequence=125 recid=20 stamp=565872502
input archive log thread=1 sequence=126 recid=21 stamp=565873739
input archive log thread=1 sequence=127 recid=24 stamp=565874852
input archive log thread=2 sequence=16 recid=15 stamp=565811247
input archive log thread=2 sequence=17 recid=18 stamp=565811458
input archive log thread=2 sequence=18 recid=19 stamp=565872502
input archive log thread=2 sequence=19 recid=22 stamp=565873742
input archive log thread=2 sequence=20 recid=23 stamp=565874800
channel ORA_DISK_1: starting piece 1 at 05/08/09
channel ORA_DISK_1: finished piece 1 at 05/08/09
piece handle=/backup/online/tpcc/0egrl456_1_1 comment=NONE
channel ORA_DISK_1: backup set complete, elapsed time: 00:00:36
Finished backup at 05/08/09

Starting Control File Autobackup at 05/08/09
piece handle=/backup/online/tpcc/c-3073235231-20050809-00 comment=NONE
Finished Control File Autobackup at 05/08/09

RMAN>

```

7-3-4. Recovery Procedure Usage Example

When a datafile diskgroup is broken.

**Shutdown an instance of database with the ABORT option.
Perform this command on OS command line of pw650-1.**

```

[pw650-1] crs_stat -t
Name                Type           Target    State    Host
-----
ora...SM1.asm       application    ONLINE    ONLINE   pw650-1
ora...-1.lsnr       application    ONLINE    ONLINE   pw650-1
ora...0-1.gsd       application    ONLINE    ONLINE   pw650-1
ora...0-1.ons       application    ONLINE    ONLINE   pw650-1
ora...0-1.vip       application    ONLINE    ONLINE   pw650-1
ora...SM2.asm       application    ONLINE    ONLINE   pw650-2
ora...-2.lsnr       application    ONLINE    ONLINE   pw650-2
ora...0-2.gsd       application    ONLINE    ONLINE   pw650-2
ora...0-2.ons       application    ONLINE    ONLINE   pw650-2
ora...0-2.vip       application    ONLINE    ONLINE   pw650-2
ora.tpcc.db         application    ONLINE    ONLINE   pw650-2
ora...c1.inst       application    ONLINE    ONLINE   pw650-1
ora...c2.inst       application    ONLINE    ONLINE   pw650-2
[pw650-1] srvctl stop database -d tpcc -o abort
[pw650-1] crs_stat -t
Name                Type           Target    State    Host
-----
ora...SM1.asm       application    ONLINE    ONLINE   pw650-1
ora...-1.lsnr       application    ONLINE    ONLINE   pw650-1
ora...0-1.gsd       application    ONLINE    ONLINE   pw650-1
ora...0-1.ons       application    ONLINE    ONLINE   pw650-1
ora...0-1.vip       application    ONLINE    ONLINE   pw650-1
ora...SM2.asm       application    ONLINE    ONLINE   pw650-2
ora...-2.lsnr       application    ONLINE    ONLINE   pw650-2
ora...0-2.gsd       application    ONLINE    ONLINE   pw650-2
ora...0-2.ons       application    ONLINE    ONLINE   pw650-2
ora...0-2.vip       application    ONLINE    ONLINE   pw650-2
ora.tpcc.db         application    OFFLINE   OFFLINE
ora...c1.inst       application    OFFLINE   OFFLINE
ora...c2.inst       application    OFFLINE   OFFLINE
[pw650-1]

```

**Dismount a disk group of datafile.
Perform this command on ASM instance of pw650-1.**

SQL> select group_number, name, state, type from v\$asm_diskgroup;

GROUP_NUMBER	NAME	STATE	TYPE
1	DG_ARCH	MOUNTED	EXTERN
2	DG_DATA	MOUNTED	EXTERN
3	DG_REDO	MOUNTED	EXTERN

SQL> alter diskgroup DG_DATA dismount;

Diskgroup altered.

SQL> select group_number, name, state, type from v\$asm_diskgroup;

GROUP_NUMBER	NAME	STATE	TYPE
1	DG_ARCH	MOUNTED	EXTERN
0	DG_DATA	DISMOUNTED	
3	DG_REDO	MOUNTED	EXTERN

SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2 from v\$asm_disk ad, v\$asm_diskgroup ag
3 where ad.group_number != 0 and ad.group_number = ag.group_number;

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
	CACHED	CANDIDATE	/dev/sfdsk/tpcc/rdisk/data1
	CLOSED	MEMBER	/dev/sfdsk/tpcc/rdisk/data2
DG_REDO	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/redo
DG_ARCH	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/arch

SQL>

**Dismount a disk group of datafile.
Perform this command on ASM instance of pw650-2.**

SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2 from v\$asm_disk ad, v\$asm_diskgroup ag
3 where ad.group_number != 0 and ad.group_number = ag.group_number;

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
DG_REDO	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/redo
DG_DATA	CACHED	CANDIDATE	/dev/sfdsk/tpcc/rdisk/data1
DG_DATA	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/data2
DG_ARCH	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/arch

SQL> select group_number, name, state, type from v\$asm_diskgroup;

GROUP_NUMBER	NAME	STATE	TYPE
1	DG_ARCH	MOUNTED	EXTERN
2	DG_DATA	MOUNTED	EXTERN
3	DG_REDO	MOUNTED	EXTERN

SQL> alter diskgroup DG_DATA dismount;

Diskgroup altered.

SQL> select group_number, name, state, type from v\$asm_diskgroup;

GROUP_NUMBER	NAME	STATE	TYPE
1	DG_ARCH	MOUNTED	EXTERN
0	DG_DATA	DISMOUNTED	
3	DG_REDO	MOUNTED	EXTERN

SQL>

Restore all disks belonging to a disk group of datafile using Advanced Copy feature.

```
# /opt/FJVSswsts/bin/swstrestore /dev/sfdsk/tpcc/dsk/data1:mplb2052
/dev/sfdsk/tpcc/dsk/data1:mplb2052 swstrestore completed
# /opt/FJVSswsts/bin/swstrestore /dev/sfdsk/tpcc/dsk/data2:mplb2054, ad.path
/dev/sfdsk/tpcc/dsk/data2:mplb2054 swstrestore completed
# 3 where ad.group_number != 0 and ad.group_number = ag.group_number;
```

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
	CLOSED	MEMBER	/dev/sfdsk/tpcc/rdisk/data1
	CLOSED	MEMBER	/dev/sfdsk/tpcc/rdisk/data2
DG_REDO	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/redo
DG_ARCH	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/arch

SQL> select group_number, name, state, type from v\$asm_diskgroup;

GROUP_NUMBER	NAME	STATE	TYPE
1	DG_ARCH	MOUNTED	EXTERN
0	DG_DATA	DISMOUNTED	
3	DG_REDO	MOUNTED	EXTERN

SQL> alter diskgroup DG_DATA mount;

Diskgroup altered.

```
SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2 from v$asm_disk ad, v$asm_diskgroup ag
3 where ad.group_number != 0 and ad.group_number = ag.group_number;
```

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
DG_REDO	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/redo
DG_DATA	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/data1
DG_DATA	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/data2
DG_ARCH	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/arch

SQL> select group_number, name, state, type from v\$asm_diskgroup;

GROUP_NUMBER	NAME	STATE	TYPE
1	DG_ARCH	MOUNTED	EXTERN
2	DG_DATA	MOUNTED	EXTERN
3	DG_REDO	MOUNTED	EXTERN

SQL>

**Startup an instance of database with the MOUNT option.
Perform this command on recovery manager of pw650-1.**

```

RMAN> startup mount

Oracle instance started
database mounted

Total System Global Area      1610612736 bytes

Fixed Size                    1302736 bytes
Variable Size                 399253296 bytes
Database Buffers              1207959552 bytes
Redo Buffers                   2097152 bytes

RMAN>

```

**Perform media recovery with a complete recovery.
Perform this command on recovery manager of pw650-1.**

```

RMAN> recover database;

Starting recover at 05/08/09
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=152 devtype=DISK

starting media recovery
media recovery complete

Finished recover at 05/08/09

RMAN>

```

**Open the database.
Perform this command on recovery manager of pw650-1.**

```

RMAN> alter database open;

database opened

RMAN>

```

When all diskgroups are broken.

**Shutdown an instance of database with the ABORT option.
Perform this command on OS command line of pw650-1.**

```

[pw650-1] crs_stat -t

```

Name	Type	Target	State	Host
ora....SM1.asm	application	ONLINE	ONLINE	pw650-1
ora....-1.lsnr	application	ONLINE	ONLINE	pw650-1
ora....0-1.gsd	application	ONLINE	ONLINE	pw650-1
ora....0-1.ons	application	ONLINE	ONLINE	pw650-1
ora....0-1.vip	application	ONLINE	ONLINE	pw650-1
ora....SM2.asm	application	ONLINE	ONLINE	pw650-2
ora....-2.lsnr	application	ONLINE	ONLINE	pw650-2
ora....0-2.gsd	application	ONLINE	ONLINE	pw650-2
ora....0-2.ons	application	ONLINE	ONLINE	pw650-2
ora....0-2.vip	application	ONLINE	ONLINE	pw650-2
ora.tpcc.db	application	ONLINE	ONLINE	pw650-2
ora....c1.inst	application	ONLINE	ONLINE	pw650-1

```
ora...c2.inst application ONLINE ONLINE pw650-2
[pw650-1] srvctl stop database -d tpcc -o abort
[pw650-1] crs_stat -t
Name                Type                Target State      Host
-----
ora...SM1.asm       application         ONLINE ONLINE    pw650-1
ora...-1.lsnr       application         ONLINE ONLINE    pw650-1
ora...0-1.gsd       application         ONLINE ONLINE    pw650-1
ora...0-1.ons       application         ONLINE ONLINE    pw650-1
ora...0-1.vip       application         ONLINE ONLINE    pw650-1
ora...SM2.asm       application         ONLINE ONLINE    pw650-2
ora...-2.lsnr       application         ONLINE ONLINE    pw650-2
ora...0-2.gsd       application         ONLINE ONLINE    pw650-2
ora...0-2.ons       application         ONLINE ONLINE    pw650-2
ora...0-2.vip       application         ONLINE ONLINE    pw650-2
ora.tpcc.db         application         OFFLINE OFFLINE
ora...c1.inst       application         OFFLINE OFFLINE
ora...c2.inst       application         OFFLINE OFFLINE
[pw650-1]
```

**Dismount a disk group of datafile.
Perform this command on ASM instance of pw650-1.**

```
SQL> alter diskgroup DG_DATA dismount;
```

Diskgroup altered.

```
SQL>
```

**Dismount a disk group of datafile, a disk group of redo log and a disk group of archive log.
Perform this command on ASM instance of pw650-2.**

```
SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2 from v$asm_disk ad, v$asm_diskgroup ag
3 where ad.group_number != 0 and ad.group_number = ag.group_number;
```

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
DG_REDO	CACHED	CANDIDATE	/dev/sfdsk/tpcc/rdisk/redo
DG_DATA	CACHED	CANDIDATE	/dev/sfdsk/tpcc/rdisk/data1
DG_DATA	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/data2
DG_ARCH	CACHED	CANDIDATE	/dev/sfdsk/tpcc/rdisk/arch

```
SQL> select group_number, name, state, type from v$asm_diskgroup;
```

GROUP_NUMBER	NAME	STATE	TYPE
1	DG_ARCH	MOUNTED	EXTERN
2	DG_DATA	MOUNTED	EXTERN
3	DG_REDO	MOUNTED	EXTERN

```
SQL> alter diskgroup DG_DATA, DG_REDO, DG_ARCH dismount;
```

Diskgroup altered.

```
SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2 from v$asm_disk ad, v$asm_diskgroup ag
3 where ad.group_number != 0 and ad.group_number = ag.group_number;
```

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
	CLOSED	CANDIDATE	/dev/sfdsk/tpcc/rdisk/arch
	CLOSED	CANDIDATE	/dev/sfdsk/tpcc/rdisk/data1
	CLOSED	MEMBER	/dev/sfdsk/tpcc/rdisk/data2

```
CLOSED    CANDIDATE    /dev/sfdsk/tpcc/rdisk/redo
```

```
SQL> select group_number, name, state, type from v$asm_diskgroup;
```

```
GROUP_NUMBER NAME          STATE          TYPE
-----
0    DG_DATA          DISMOUNTED
```

```
SQL>
```

**Drop the other disk groups.
Perform this command on ASM instance of pw650-1.**

```
SQL> drop diskgroup DG_REDO including contents;
```

```
Diskgroup dropped.
```

```
SQL> drop diskgroup DG_ARCH including contents;
```

```
Diskgroup dropped.
```

```
SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2 from v$asm_disk ad, v$asm_diskgroup ag
3 where ad.group_number != 0 and ad.group_number = ag.group_number;
```

```
GROUP_NAME    MOUNT_S    HEADER_STATU    PATH
-----
                CLOSED    CANDIDATE    /dev/sfdsk/tpcc/rdisk/arch
                CLOSED    CANDIDATE    /dev/sfdsk/tpcc/rdisk/data1
                CLOSED    MEMBER      /dev/sfdsk/tpcc/rdisk/data2
                CLOSED    CANDIDATE    /dev/sfdsk/tpcc/rdisk/redo
```

```
SQL> select group_number, name, state, type from v$asm_diskgroup;
```

```
GROUP_NUMBER NAME          STATE          TYPE
-----
0    DG_DATA          DISMOUNTED
```

```
SQL>
```

**Initialize disk header of all disks belong to DG_REDO, and
DG_ARCH.
Perform this command on ASM instance of pw650-1.**

```
[pw650-1] dd if=/dev/zero of=/dev/sfdsk/tpcc/rdisk/redo bs=1024k count=10
10+0 records in
10+0 records out
[pw650-1] dd if=/dev/zero of=/dev/sfdsk/tpcc/rdisk/arch bs=1024k count=10
10+0 records in
10+0 records out
[pw650-1]
```

**Create a disk group of DG_REDO and a disk group of DG_ARCH.
Perform this command on ASM instance of pw650-1.**

```
SQL> create diskgroup DG_REDO external redundancy disk
2 '/dev/sfdsk/tpcc/rdisk/redo';
```

```
Diskgroup created.
```

```
SQL> create diskgroup DG_ARCH external redundancy disk
2 '/dev/sfdsk/tpcc/rdisk/arch';
```

```
Diskgroup created.
```

```
SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2  from v$asm_disk ad, v$asm_diskgroup ag
3  where ad.group_number != 0 and ad.group_number = ag.group_number;
```

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
	CLOSED	CANDIDATE	/dev/sfdsk/tpcc/rdisk/data1
	CLOSED	MEMBER	/dev/sfdsk/tpcc/rdisk/data2
DG_REDO	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/redo
DG_ARCH	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/arch

```
SQL> select group_number, name, state, type from v$asm_diskgroup;
```

GROUP_NUMBER	NAME	STATE	TYPE
0	DG_DATA	DISMOUNTED	
1	DG_ARCH	MOUNTED	EXTERN
2	DG_REDO	MOUNTED	EXTERN

```
SQL>
```

**Restore all disks belonging to a disk group of datafile using Advanced Copy feature.
Perform this command on OS command line of pw650-1.**

```
# /opt/FJSVswsts/bin/swstrestore /dev/sfdsk/tpcc/dsk/data1:mplb2052
/dev/sfdsk/tpcc/dsk/data1:mplb2052 swstrestore completed
# /opt/FJSVswsts/bin/swstrestore /dev/sfdsk/tpcc/dsk/data2:mplb2054
/dev/sfdsk/tpcc/dsk/data2:mplb2054 swstrestore completed
#
```

**Mount a disk group of datafile.
Perform this command on ASM instance of pw650-1.**

```
SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2  from v$asm_disk ad, v$asm_diskgroup ag
3  where ad.group_number != 0 and ad.group_number = ag.group_number;
```

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
	CLOSED	MEMBER	/dev/sfdsk/tpcc/rdisk/data1
	CLOSED	MEMBER	/dev/sfdsk/tpcc/rdisk/data2
DG_REDO	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/redo
DG_ARCH	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/arch

```
SQL> select group_number, name, state, type from v$asm_diskgroup;
```

GROUP_NUMBER	NAME	STATE	TYPE
0	DG_DATA	DISMOUNTED	
1	DG_ARCH	MOUNTED	EXTERN
2	DG_REDO	MOUNTED	EXTERN

```
SQL> alter diskgroup DG_DATA mount;
```

Diskgroup altered.

```
SQL> select ag.name "GROUP_NAME", ad.mount_status, ad.header_status, ad.path
2  from v$asm_disk ad, v$asm_diskgroup ag
3  where ad.group_number != 0 and ad.group_number = ag.group_number;
```

GROUP_NAME	MOUNT_S	HEADER_STATU	PATH
DG_REDO	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/redo
DG_DATA	CACHED	MEMBER	/dev/sfdsk/tpcc/rdisk/data1

```
DG_DATA      CACHED    MEMBER      /dev/sfdsk/tpcc/rdisk/data2
DG_ARCH      CACHED    MEMBER      /dev/sfdsk/tpcc/rdisk/arch
```

```
SQL> select group_number, name, state, type from v$asm_diskgroup;
```

```
GROUP_NUMBER NAME          STATE        TYPE
-----
3 DG_DATA      MOUNTED     EXTERN
1 DG_ARCH      MOUNTED     EXTERN
2 DG_REDO      MOUNTED     EXTERN
```

```
SQL>
```

**Copy a backup control file to work area.
Perform this command on OS command line of pw650-1.**

```
[pw650-1] cd /backup/online/tpcc
[pw650-1] cp -p c-3073235231-20050812-00 /backup/ctl00.bak
[pw650-1]
```

**Startup an instance of database without mount.
Perform this command on recovery manager of pw650-1.**

```
RMAN> startup nomount;
```

```
Oracle instance started
```

```
Total System Global Area      1610612736 bytes
```

```
Fixed Size                      1302736 bytes
Variable Size                    399253296 bytes
Database Buffers                 1207959552 bytes
Redo Buffers                      2097152 bytes
```

```
RMAN>
```

**Restore the control file from a control file at work area.
Perform this command on recovery manager of pw650-1.**

```
RMAN> restore controlfile from '/backup/ctl00.bak';
```

```
Starting restore at 2005-08-12 13:49:09
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=151 devtype=DISK
```

```
channel ORA_DISK_1: restoring controlfile
channel ORA_DISK_1: restore complete
output filename=+DG_REDO/tpcc/tpcc_cnt1
output filename=+DG_ARCH/tpcc/tpcc_cnt2
Finished restore at 2005-08-12 13:49:20
```

```
RMAN>
```

**Mount the database.
Perform this command on recovery manager of pw650-1.**

```
RMAN> alter database mount;
```

```
database mounted
released channel: ORA_DISK_1
```

```
RMAN>
```

**Perform media recovery with an incomplete recovery.
Perform this command on recovery manager of pw650-1.**

```

RMAN> recover database until sequence 127 thread 1;

```

```

Starting recover at 2005-08-12 13:51:47
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=151 devtype=DISK

```

```

starting media recovery

```

```

channel ORA_DISK_1: starting archive log restore to default destination
channel ORA_DISK_1: restoring archive log
archive log thread=1 sequence=126
channel ORA_DISK_1: restoring archive log
archive log thread=2 sequence=19
channel ORA_DISK_1: restored backup piece 1
piece handle=/backup/online/tpcc/0eprt87f_1_1 tag=TAG20050812T132607
channel ORA_DISK_1: restore complete
archive log filename=+DG_ARCH/tpcc/1_126_565531999.arc thread=1 sequence=126
archive log filename=+DG_ARCH/tpcc/2_19_565531999.arc thread=2 sequence=19
media recovery complete
Finished recover at 2005-08-12 13:53:37

```

```

RMAN>

```

**Open the database with the RESETLOGS option.
Perform this command on recovery manager of pw650-1.**

```

RMAN> alter database open resetlogs;

```

```

database opened
new incarnation of database registered in recovery catalog
starting full resync of recovery catalog
full resync complete

```

```

RMAN>

```

**Add any temporary files to a temporary tablespace.
Perform this command on database instance of pw650-1.**

```

SQL> select ts.tablespace_name, file_name from v$tempfile;

```

```

no rows selected

```

```

SQL> alter tablespace TEMP add tempfile
  2 '+DG_DATA/TPCC/TEMP1' size 3784M reuse,
  3 '+DG_DATA/TPCC/TEMP2' size 3784M reuse;

```

```

Tablespace altered.

```

```

SQL> select ts.tablespace_name, file_name from v$tempfile;

```

```

TABLESPACE  FILE_NAME
-----
TEMP        +DG_DATA/tpcc/temp1
TEMP        +DG_DATA/tpcc/temp2

```

```

SQL>

```

**Delete some invalid entries of archive log files.
Perform this command on recovery manager of pw650-1.**

```

RMAN> crosscheck archivelog all;

```

```

allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=134 devtype=DISK

```

```
validation failed for archived log
archive log filename=+DG_ARCH/tpcc/1_123_565531999.arc recid=16 stamp=565811250
validation failed for archived log
archive log filename=+DG_ARCH/tpcc/1_124_565531999.arc recid=17 stamp=565811456
validation failed for archived log
archive log filename=+DG_ARCH/tpcc/1_125_565531999.arc recid=20 stamp=566140624
validation succeeded for archived log
archive log filename=+DG_ARCH/tpcc/1_126_565531999.arc recid=24 stamp=566142746
validation failed for archived log
archive log filename=+DG_ARCH/tpcc/2_16_565531999.arc recid=15 stamp=565811247
validation failed for archived log
archive log filename=+DG_ARCH/tpcc/2_17_565531999.arc recid=18 stamp=565811458
validation failed for archived log
archive log filename=+DG_ARCH/tpcc/2_18_565531999.arc recid=19 stamp=566140623
validation succeeded for archived log
archive log filename=+DG_ARCH/tpcc/2_19_565531999.arc recid=23 stamp=566142717
Crosschecked 8 objects
```

```
RMAN> list archivelog all;
```

List of Archived Log Copies

Key	Thrd	Seq	S	Low	Time	Name
501	1	123	X	2005-08-05	20:02:31	+DG_ARCH/tpcc/1_123_565531999.arc
510	1	124	X	2005-08-08	17:47:28	+DG_ARCH/tpcc/1_124_565531999.arc
665	1	125	X	2005-08-08	17:50:56	+DG_ARCH/tpcc/1_125_565531999.arc
1758	1	126	A	2005-08-12	13:16:59	+DG_ARCH/tpcc/1_126_565531999.arc
500	2	16	X	2005-08-05	20:02:32	+DG_ARCH/tpcc/2_16_565531999.arc
511	2	17	X	2005-08-08	17:47:26	+DG_ARCH/tpcc/2_17_565531999.arc
664	2	18	X	2005-08-08	17:50:57	+DG_ARCH/tpcc/2_18_565531999.arc
1757	2	19	A	2005-08-12	13:16:58	+DG_ARCH/tpcc/2_19_565531999.arc

```
RMAN> delete expired archivelog all;
```

```
released channel: ORA_DISK_1
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=134 devtype=DISK
```

List of Archived Log Copies

Key	Thrd	Seq	S	Low	Time	Name
501	1	123	X	2005-08-05	20:02:31	+DG_ARCH/tpcc/1_123_565531999.arc
510	1	124	X	2005-08-08	17:47:28	+DG_ARCH/tpcc/1_124_565531999.arc
665	1	125	X	2005-08-08	17:50:56	+DG_ARCH/tpcc/1_125_565531999.arc
500	2	16	X	2005-08-05	20:02:32	+DG_ARCH/tpcc/2_16_565531999.arc
511	2	17	X	2005-08-08	17:47:26	+DG_ARCH/tpcc/2_17_565531999.arc
664	2	18	X	2005-08-08	17:50:57	+DG_ARCH/tpcc/2_18_565531999.arc

```
Do you really want to delete the above objects (enter YES or NO)? y
```

```
deleted archive log
archive log filename=+DG_ARCH/tpcc/1_123_565531999.arc recid=16 stamp=565811250
deleted archive log
archive log filename=+DG_ARCH/tpcc/1_124_565531999.arc recid=17 stamp=565811456
deleted archive log
archive log filename=+DG_ARCH/tpcc/1_125_565531999.arc recid=20 stamp=566140624
deleted archive log
archive log filename=+DG_ARCH/tpcc/2_16_565531999.arc recid=15 stamp=565811247
deleted archive log
archive log filename=+DG_ARCH/tpcc/2_17_565531999.arc recid=18 stamp=565811458
deleted archive log
archive log filename=+DG_ARCH/tpcc/2_18_565531999.arc recid=19 stamp=566140623
Deleted 6 EXPIRED objects
```

```
RMAN> list archivelog all;
```

```
List of Archived Log Copies
```

Key	Thrd	Seq	S	Low Time	Name
1758	1	126	A	2005-08-12 13:16:59	+DG_ARCH/tpcc/1_126_565531999.arc
1757	2	19	A	2005-08-12 13:16:58	+DG_ARCH/tpcc/2_19_565531999.arc

```
RMAN>
```

7-4. Additional Technical References

Oracle web site

- (1) Oracle Database 10g Automatic Storage Management Technical Best Practices, Nitin Vengurlekar
http://www.oracle.com/technology/products/database/asm/pdf/asm_bestpractices_9_7.pdf
- (2) Automatic Storage Management
<http://www.oracle.com/technology/products/database/asm/index.html>
- (3) Oracle Real Application Clusters
<http://www.oracle.com/technology/products/database/clustering/index.html>

Fujitsu's ETERNUS storage systems web site

- (1) Fujitsu Disk Storage Systems
<http://www.fujitsu.com/global/services/computing/storage/system/>
- (2) White Papers
<http://www.fujitsu.com/global/support/computing/storage/system/>

About This White Paper

This white paper is devoted to providing technical information and an overview of the usage of Oracle Database 10g Automatic Storage Management function and Fujitsu ETERNUS storage systems.

The contents of this document may be modified without any prior notice. Please contact FUJITSU LIMITED if you find any error in the descriptions. FUJITSU LIMITED is not responsible for any damage or loss that might be caused by the contents of this document.

FUJITSU LIMITED

Global Marketing
Platform Business Group
4-1-1, Kamikodanaka, Nakahara-ku, Kawasaki 211-8588 Japan

E-MAIL: storage-system@fujitsu.com
"Fujitsu Storage System Web Site" <http://www.fujitsu.com/storage/>