

**HITACHI**  
Inspire the Next

**ORACLE®**

# Guidelines for Optimized Backup and Recovery of Oracle 11g Database with Automatic Storage Management Using Hitachi Disk Array Replication

Technical Brief

*By Satoshi Saito, SAN Solution Division, Hitachi, Ltd.,  
Todd Hanson, Hitachi Data Systems, and  
Nitin Vengurlekar, Oracle Corporation*

October 2008



## Executive Summary

In today's competitive business environment, an irrecoverable data loss in an enterprise can quickly result in revenue loss, reputation loss and customer dissatisfaction. Enterprises are therefore looking for high quality cost-effective backup solutions to protect their business against this risk. In order to meet this challenge, Hitachi Ltd., and Oracle have partnered to deliver an industry leading backup and recovery system by combining Hitachi storage and Oracle 11g Database with Automatic Storage Management (ASM). This document is a joint effort by Hitachi and Oracle to provide guidelines for optimizing backup and recovery of Oracle 11g Database with ASM using Hitachi storage.

The Hitachi Universal Storage Platform™ V and the Hitachi Adaptable Modular Storage 2000 family provide advanced business continuity for enterprises using Oracle 11g Database by combining data protection, disaster recovery and replication functions on the business site. The platforms deliver higher value through simplified storage integration and storage management. The Universal Storage Platform V and the Adaptable Modular Storage 2000 family also enhance the performance and automation of the Oracle 11g Database with ASM.

First introduced in Oracle 10g, ASM improves the efficiency, performance and availability of the storage underlying the Oracle Database. ASM eliminates the need for third-party cluster volume management and file system functions for Oracle Databases. Oracle Recovery Manager (RMAN) offers an efficient method for backup and recovery of Oracle Databases. It also provides block level corruption detection during backup and restore.

This document examines the backup and recovery operations in an ASM environment using Hitachi replication technology. It provides guidelines for the most practical configurations. Backup and recovery in an ASM environment can be managed in the same way as they are in the traditional file system. However, it involves distinctive procedures and setting methods that are described in this document for your reference.

**For best results use Acrobat Reader 8.0**

# Contents

<b>Introduction</b> .....	<b>1</b>
New Functions in the Hitachi Adaptable Modular Storage 2000 Family .....	2
<b>Using Oracle ASM with Hitachi Storage Solutions</b> .....	<b>3</b>
Storage Settings (Settings for a Shared Volume) .....	3
Overview of Hitachi Dynamic Link Manager Software .....	3
Overview of Hitachi ShadowImage Software .....	4
Overview of Hitachi TrueCopy® and Hitachi Universal Replicator Software .....	4
Consistency Groups and At-time Split Options for Point-in-time Copies.....	4
ASM Disk Groups and Pair Groups (Consistency Groups).....	5
Basic ASM Operations.....	5
<b>Settings for ShadowImage Software</b> .....	<b>6</b>
Pair Groups .....	6
Pair Group Operations.....	6
Description of RAID Manager Commands.....	8
<b>ASM Backup/Recovery Using ShadowImage Software</b> .....	<b>8</b>
Overview.....	8
Hot Backup .....	8
Cold Backup .....	9
Backup Requirements .....	10
Recovery .....	11
Scenario 1: Hot Backup of Oracle 11g Database with ASM Using Hitachi ShadowImage Software and Oracle RMAN .....	12
Scenario 1b: Cold Backup of Oracle 11g Database with ASM Using Hitachi ShadowImage Software and Oracle RMAN .....	15
Scenario 2: Recovery of Oracle 11g Database with ASM Using Hitachi ShadowImage Software and Oracle RMAN .....	16
Scenario 3: Cloning of Oracle 11g Database with ASM Using Hitachi ShadowImage Software and Oracle RMAN .....	18
<b>Conclusion</b> .....	<b>20</b>
<b>Schedule A: RMAN Recovery Catalog Guidelines</b> .....	<b>21</b>
<b>Schedule B: Example of the TNSNAMES.ORA Configuration File</b> .....	<b>22</b>
<b>Schedule C: Example of Hitachi ShadowImage Configuration Files</b> .....	<b>23</b>
<b>Schedule D: Example of an Initialization Parameter File</b> .....	<b>24</b>

# Guidelines for Optimized Backup and Recovery of Oracle 11g Database with Automatic Storage Management Using Hitachi Disk Array Replication

Technical Brief

*By Satoshi Saito, SAN Solution Division, Hitachi, Ltd.,  
Todd Hanson, Hitachi Data Systems, and  
Nitin Vengurlekar, Oracle Corporation*


## Introduction

A Hitachi storage system provides a high performance and reliable volume replication function that can generate a point-in-time copy swiftly. This document describes the use of Hitachi ShadowImage® Heterogeneous Replication software on the Hitachi Universal Storage Platform™ V and Hitachi ShadowImage Replication software on the Hitachi Adaptable Modular Storage 2000 family for the following operations in an Oracle 11g Database with Automatic Storage Management (ASM) environment.

- Database backup
- Database recovery
- Database cloning (batch processing, migration inspection, development)

Replication using Hitachi storage systems is more efficient than a server-based solution because all the processing takes place inside the storage without any transfer of data from storage to a server. Storage-based replication also promotes centralized administration and lowers total cost of ownership (TCO) for data protection operations. Hitachi replication products use consistency group technology to meet the Oracle Database 11g ASM split-mirror guideline conditions. This enables the creation of a valid consistent copy of an ASM database even when it is online.

- Point-in-time split copies can be created in all logical units (LUs) included in each ASM disk group, and all LUs are split at the same time.
- All LUs included in each ASM disk group are atomically split while preserving their write complete status, and write ordering is preserved.
- The data table area and the archive logs are each in different ASM disk groups. (because the replications are created at different times and we wish to manage them separately)
  - The DATADG disk group includes data table areas
  - The ARCHDG disk group includes archive logs and control files
  - The REDODG disk group includes the REDO log file

- 
- All volumes of the ASM disk groups are included in the pair group of the Hitachi replication function.
  - The backup server has the same operating system as the business server, and has suitable Oracle binary files and initialized parameter files, password file and user ID/group.
  - The backup server has been configured to connect to the recovery catalog of Oracle Recovery Manager (RMAN).

An optimal backup and recovery system is achieved by using the industry's best platforms, the Hitachi Universal Storage Platform V and Hitachi Adaptable Modular Storage 2000 family, for deploying Oracle 11g Database with ASM. The advanced volume products in these Hitachi storage systems complement ASM perfectly by optimizing performance and enabling seamless backup. The combination of ASM with Hitachi storage also provides robust and reliable databases, steady backup and disk recovery, and powerful administration tools.

This document introduces the guidelines for backup and recovery when combining Oracle 11g Database with either a Hitachi Universal Storage Platform or Hitachi Adaptable Modular Storage 2000 system. It highlights the procedural changes required for benefitting from these powerful combinations.

For the purpose of convenience, this document refers to the latest models, the Hitachi Universal Storage Platform™ V and the Hitachi Adaptable Modular Storage 2000. However you can perform backup, recovery and cloning by following the same procedures with any of the following products, as their architecture and usability are the same.

- Hitachi Adaptable Modular Storage 2000 family
- Hitachi Universal Storage Platform V
- Hitachi Universal Storage Platform VM
- Hitachi Adaptable Modular Storage family
- Hitachi Workgroup Modular Storage
- Hitachi Universal Storage Platform
- Hitachi Network Storage Controller

This document also refers to ShadowImage software when discussing the volume replication function, but the function applies to both ShadowImage and the products shown below:

- Hitachi ShadowImage In-System Replication
- Hitachi ShadowImage Heterogeneous In-System Replication software
- Hitachi TrueCopy Synchronous
- Hitachi TrueCopy Extended Distance
- Hitachi TrueCopy Remote Replication software
- Hitachi Universal Replicator

## New Functions in the Hitachi Adaptable Modular Storage 2000 Family

The Hitachi Adaptable Modular Storage 2000 family includes several new features that improve storage performance and simplify storage configuration and management.

### *High-performance Controller with Load Balancing*

Data processing can be performed quickly and flexibly between controllers by using the cross path mechanism between controllers. This cross path operation minimizes overheads and enables excellent active-active operation from the host.

The automatic load balancing function detects workload imbalance between the two controllers and makes the necessary adjustments to balance data processing. This optimizes the system performance.

Cross path processing between controllers also enables online maintenance (upgrade) of controller firmware without having to use path switching software.

### *Enhanced Performance Using SAS (Serial Attached SCSI) I/F Technology*

SAS is a serial interface that provides highly reliable point-to-point architecture. Unlike complicated loop topology, it links devices on a 1:1 basis in a network via Expander.

SAS provides six times larger bandwidth for connection as compared to cables which were used for connecting devices in the previous Adaptable Modular Storage systems. In addition, SAS delivers higher performance and expanded bandwidth by implementing a WL (wide link). SAS connects controllers with expanders by bundling four paths (links).

The physical interface of SAS is also compatible with a SATA interface. This allows a SATA hard disk drive to be supported in the same enclosure with a native protocol, without the need for a mechanism to convert from the original back-end Fiber Channel to SATA.

### *Improved Usability with Storage Navigator Modular 2*

Storage Navigator Modular 2 has a wizard-based GUI interface with online help. It enables easy storage installation and administration, without having to refer to manuals and other reference materials.

It makes the copy function (snapshot and internal volume replication functions) easier to use by providing GUI-based settings.

## Using Oracle ASM with Hitachi Storage Solutions

### Storage Settings (Settings for a Shared Volume)

A shared volume is required to deploy Oracle Database with Real Application Cluster (RAC). In order to create a shared volume, multiple servers should be able to recognize the volume. This can be achieved through configuration of the operating system (OS) platform. Shared volume may also be managed using Hitachi Dynamic Link Manager software, which supports path redundancy and is used to define multiple paths from servers to a volume. An overview of Dynamic Link Manager is provided in the following section.

### Overview of Hitachi Dynamic Link Manager Software

Hitachi Dynamic Link Manager manages the access paths to storage systems. It enhances the system's availability and reliability by providing a function that ensures failover of an access path to another path in the event of a fault on the path.

Hitachi Dynamic Link Manager provides the following features:

- **Load balancing.** When multiple access paths connect a host to a storage system, Dynamic Link Manager distributes the load over all the access paths. This avoids data access congestion on a specific access path.

- **Fault tolerance.** If a fault occurs on an access path when multiple access paths connect a host to a storage system, Dynamic Link Manager automatically switches to another access path. This allows the process to continue without being affected by the fault.
- **Automatic failback.** This automatically restores an access path online when it has recovered from a fault. Usually the administrator restores a failed access path at a convenient time that minimizes the impact on business. This means that failback is usually performed manually. However, it is also possible to perform failback automatically. When "automatic failback" is activated, Dynamic Link Manager automatically activates the path after the user has solved the physical route problem.
- **Automatic path health check.** Dynamic Link Manager detects faults by checking the status of access paths at intervals defined by the user. This allows path errors to be identified and corrected.

In Linux, device names such as /dev/sda and /dev/sdb are assigned to the paths to the LUN. Dynamic Link Manager finds the path to the same LUN, and assigns a shared name such as /dev/sddlm aa. The Dynamic Link Manager device format in Linux is /dev/sddlm [a-p][a-p][1-15]. The device partition is indicated by the number after sddlm\* name. All devices are displayed without a specified number. You must use the same Dynamic Link Manager device name that is used by the load balancing and failover functions. For more information, please see the product's user guide for each OS platform.

## Overview of Hitachi ShadowImage Software

Hitachi ShadowImage Heterogeneous Replication and ShadowImage Replication software provide a storage-based volume replication function that can be used in a RAID volume inside a storage system. It can also be used for external storage volumes that are virtualized using Hitachi Universal Volume Manager. ShadowImage software's primary volume (P-VOL) can create multi-generational secondary volumes (S-VOL). The ShadowImage replication function can be run without stopping I/O operations on the business server. ShadowImage has an at-time split option; this creates a simultaneous split across multiple volumes for supporting backup taken whilst an application remains online. This allows a flexible backup using a consistency group.

## Overview of Hitachi TrueCopy<sup>®</sup> and Hitachi Universal Replicator Software

Hitachi TrueCopy Synchronous/Extended Distance and Universal Replicator software provide functions used for making remote copies. They allow continuous copying by sending updated information to the remote site while continuing read/write I/O on the volume in the primary site. TrueCopy software supports synchronous replication to the remote site; data written at the primary site is not committed until the remote write is also confirmed. Universal Replicator makes asynchronous copies to the secondary site, and is designed to guarantee write ordering in order to protect the data being written by an application should a fault occur. Flexible system building is made possible by the ability to preserve the write ordering for each consistency group. Universal Replicator uses an enormous journal volume in order to back up updated information. Hitachi TrueCopy and Universal Replicator provide various recovery options to suit the system operation. For example, it is possible to create a cascaded configuration by using TrueCopy or Universal Replicator in conjunction with ShadowImage; this would permit local and remote backups to be performed.

## Consistency Groups and At-time Split Options for Point-in-time Copies

Consistency groups can be created by applying the at-time split option to a ShadowImage pair group operation. When the pair split command is executed from the server to the Universal Storage Platform or Adaptable Modular Storage system via RAID Manager, all LUs in the same consistency group are simultaneously and atomically split into pairs and secondary volume (S-VOLs) data is established. The contents of the S-VOLs are the same as the P-VOL when the pair split operation was executed.

## ASM Disk Groups and Pair Groups (Consistency Groups)

The relationship between ASM disk groups and ShadowImage pair groups is shown in Table 1.

**Table 1. Layout for a Three Disk Group Configuration**

<b>ASM Disk Group</b>	<b>ShadowImage Pair Group (Consistency Group)</b>		<b>Use</b>
<b>DATADG</b>	DATADG	SIALLDG	SYSTEM table area SYSAUX table area UNDO table area USERS table area TEMP table area Server parameter file (spfile)
<b>ARCHDG</b>	ARCHDG		Archive REDO log file area (archive log file) Control file 1
<b>REDODG</b>			Online REDO log file node 1 (2 member 3 group) Online REDO log file node 2 (2 member 3 group) Control file 2

## Basic ASM Operations

The following section provides a brief overview of the commands used for basic Oracle ASM operations.

### *Starting and Stopping*

You can start and stop a single instance of Oracle ASM by setting +ASM as the SID environment variable, as you would for a normal database instance.

```
$ export ORACLE_SID=+ASM  
$ sqlplus / as sysdba  
SQL> startup
```

Stopping an ASM instance (single instance):

```
$ export ORACLE_SID=+ASM  
$ sqlplus / as sysdba  
SQL> shutdown
```

The srvctl command starts and stops an ASM instance in an Oracle RAC environment. This command can be executed from any node.

Starting an ASM instance (RAC environment):

```
$ srvctl start asm -n nodename1  
$ srvctl start asm -n nodename2
```

Stopping an ASM instance (RAC environment):

```
$ srvctl stop asm -n nodename1  
$ srvctl stop asm -n nodename2
```

## Checking the status of an ASM instance

Use the following query on v\$instance to check whether an ASM instance has been started or not.

```
SQL> select status from v$instancev  
MOUNTED
```

# Settings for ShadowImage Software

## Pair Groups

Backup in ShadowImage is performed in units called pair groups. When a point-in-time copy is made, these pair groups become consistency groups, and operations on the group occur simultaneously on all volumes in the group. You set a pair group using a configuration file (horcm.conf) or by use of Hitachi Replication Manager software.

For example, in a three disk configuration (consisting of the DATADG data table area, the ARCGDG archive area, and the REDODG REDO log file), the disk groups that need to be backed up are DATADG and ARCHDG. However, these two disk groups are required to be unified into a pair group. When adding or removing a LU to or from a disk group, the pair group settings are also changed.

## Pair Group Operations

The interfaces for operating the ShadowImage pair groups are (a) RAID Manager, a command line interface (CLI), (b) Storage Navigator/Storage Navigator Modular 2, GUI based software and (3) Hitachi Replication Manager, a solution for managing and monitoring large numbers of replication relationships across multiple storage platforms.

RAID Manager runs on a server. By configuring the server to recognize a specific Logical Unit called a command device (CMD), you can duplicate volumes and achieve pair status via the CMD. You need to start two instances in order to use RAID Manager; one manages the P-VOL, the other the S-VOL. Two configuration files are therefore required. Examples of the configuration files (horcm.conf) have been included in Schedule C at the end of this document for your reference.

Example of starting instances (starting instance 0 and 1):

```
# horcmstart.sh 0 1
```

Once configuration files are defined and the two horcm instances have been started, preparations for creating a ShadowImage pair are complete. The paircreate command can then be used to create a pair. The options for the paircreate command are given below.

The "-g" option specifies the pair group name.


The "-vl" option indicates the copy direction; local-to-remote. Reverse direction "-vr" option is supported but is rarely used.

The "-m grp" option specifies that the pair group will be created as a consistency group.

When creating a pair in an ASM environment, the pair group needs to be created as a consistency group. So, the point-in-time option "-m grp" is used with the "paircreate" command.

Example of creating a ShadowImage consistency group:

```
# paircreate -g DATADG -vl -m grp  
# paircreate -g ARCHDG -vl -m grp
```



A disk group usually contains multiple volumes (LUs). However, by converting it into a consistency group and attaching the point-in-time option to it, a consistent and simultaneous split can be taken across all the associated LUs without any time lag.

If the pair is created without attaching the point-in-time option, consistency may not be maintained during a split because each volume is split at different times. An inconsistent copy cannot be used for recovery. Furthermore, it is not possible to determine inconsistency until recovery.

Example of a point-in-time copy (consistency group):

```
# pairsplit -g DATADG  
# pairsplit -g ARCHDG
```

Commands such as pairsplit operate in the background, so you need to use the pairevwait command to check the transition into the correct pair status. The main options of the pairevwait command are shown below.

The "-g" option specifies the pair group name.

The "-s" option specifies the transition target pair status.

The "-t" option specifies the timeout period for status transition in seconds.

Example of checking the suspend status (PSUS):

```
# pairevwait -g DATADG -s psus -t 3600  
# pairevwait -g ARCHDG -s psus -t 3600
```

## Description of RAID Manager Commands

Table 2 shows the basic pair operation commands in RAID Manager and their functions.

**Table 2. Pair Command**

<b>Command Name</b>	<b>Description</b>	<b>Notes</b>
<b>Paircreate</b>	Creates a new pair.	Creates pair (SMPL→PAIR)
<b>Pairsplit</b>	Splits pairs and creates a backup in the secondary volume by stopping the update of the secondary volume. The pair relationship of the primary volume is maintained, and differential information is also maintained.	Splits pairs/suspends pairs (PAIR/COPY→PSUS)
<b>Pairresync</b>	Resynchronizes pairs from pair suspend status. Uses differential information to resynchronize at high speed and place in pair status. Only changed data needs to be copied.	Resynchronizes pairs (PSUS→PAIR)
<b>pairresync -restore</b>	Resynchronizes a pair from suspend status in the opposite direction and restores backup data. Uses differential information to restore at high speed. The pair is usually suspended immediately after being restored, preserving the data.	Restores pairs
<b>pairsplit -S</b>	Discards a pair's differential information and deletes the pair. To place it back into pair status, you need to create the pair again.	Deletes pair (PAIR/COPY/PSUS→SMPL)
<b>Pairwait</b>	Continues to wait until a group enters the specified status. Use to ensure pair status transition.	
<b>Pairdisplay</b>	Checks pair status and pair progress rate.	
<b>Pairvolchk</b>	Checks whether there is a point-in-time option (consistency group setting).	

## ASM Backup/Recovery Using ShadowImage Software

### Overview

This section provides guidelines for the backup, recovery and cloning scenarios of an Oracle 11g Database with ASM using Hitachi ShadowImage software and Oracle RMAN. The following scenarios are discussed in this section.

Scenario 1A: Hot Backup of Oracle 11g Database with ASM using Hitachi ShadowImage and Oracle RMAN

Scenario 1B: Cold Backup of Oracle 11g Database with ASM using Hitachi ShadowImage and Oracle RMAN

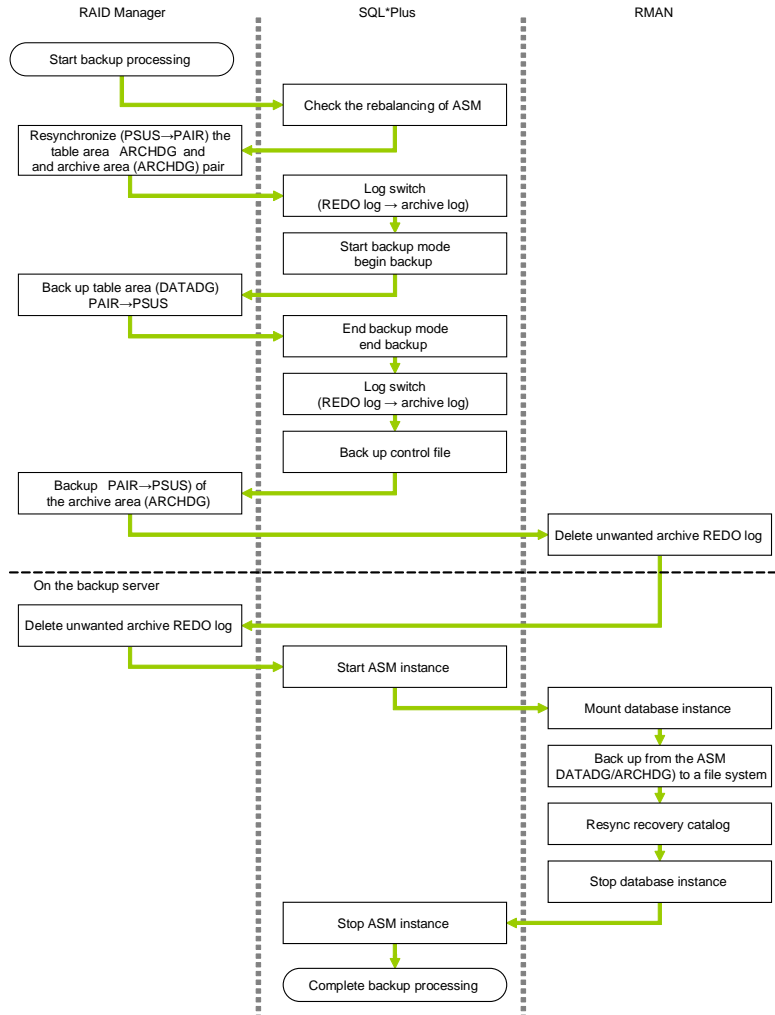
Scenario 2: Recovery of Oracle 11g Database with ASM backed up in Scenario 1A using Hitachi ShadowImage and Oracle RMAN

Scenario 3: Cloning of Oracle 11g Database with ASM using Hitachi ShadowImage and Oracle RMAN

### Hot Backup

When performing a hot backup in an ASM environment, the point-in-time option must be used during pair creation. If the database is backed up by splitting a pair without the point-in-time option, consistency between the P-VOLs and S-VOLs will not be preserved, thereby creating a risk of invalid data.

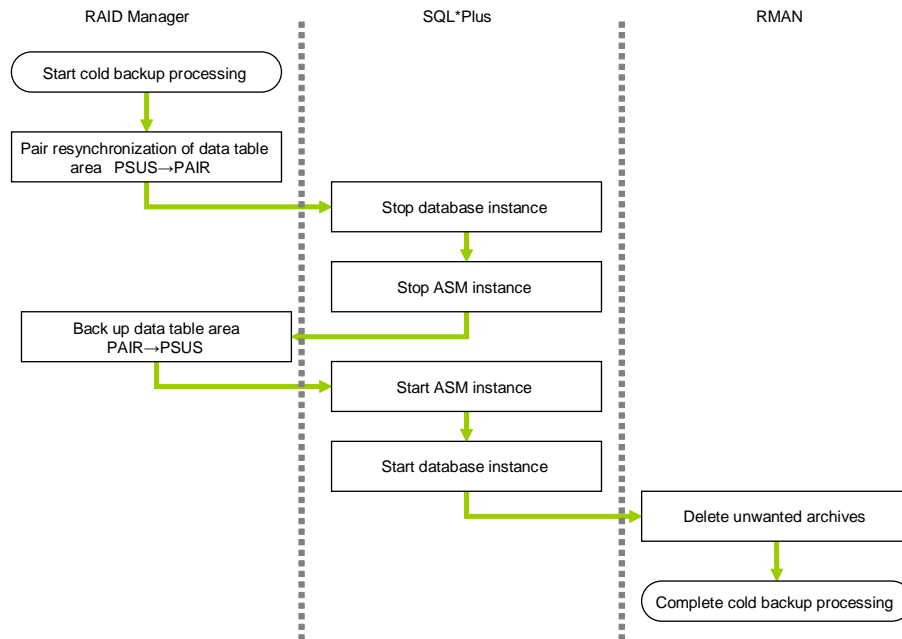
Figure 1. Flowchart for Hot Backup



## Cold Backup

A cold backup provides a consistent backup without the need for any additional recovery. To perform a cold backup of Oracle 11g Database with ASM, the database instance is stopped first followed by the termination of the ASM instance. The volume is then replicated using Hitachi ShadowImage. After the backup is completed, start the ASM instance first, and then start the database instance.

**Figure 2. Flowchart for Cold Backup**



## Backup Requirements

### *Backup of Data Files*

Oracle11g Database can be put into backup mode with just one command. Up until Oracle 9i, hot backup scripts were required to change the tablespaces one-by-one into backup mode. The single backup command for all the tablespaces simplifies the backup process in Oracle 11g Database.

Starting backup mode:

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

Ending backup mode:

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

### *Backup of the Control File*

Appropriate copies of the control files are required in many restore scenarios. You can back up the control file by using Oracle Recovery Manager (RMAN) or SQL\*Plus. The control file needs to be backed up separately since this step is done after the backup of tablespaces.

```
RMAN> copy current controlfile to '+ARCHDGC/control_file/control_bak';
```

There is also a method for creating the control file creation script. If you use "backup control file to trace," you can create a file that includes the "create control file" script. If the output destination is not specified by its full path, the file will be created in the directory specified by USER\_DUMP\_DEST.

```
SQL> alter database backup controlfile to trace as '/tmp/control.bak';
```

### *REDO Log File and Archive Log File (Archive Log Area)*

ShadowImage does not back up the REDO log file. Oracle does not recommend the backup of the REDO log file as corruption may occur, if it is restored. But REDO information is required for recovery and is backed up using the "alter system archive log current" command to copy it to the archive area.

Archive log files that would have been created on the old file system are created in a disk group in an ASM environment.

As the archive area has no backup mode, an invalid archive log may be created if a backup is done while creating an archive log file. Before backing up archive area using ShadowImage, an archive log file for each node can be created by using "alter system archive log current" command. Creation of the archive log files can be verified by looking at Dictionary View and the alert logs of each instance. Once the backup has been completed, the old archive log files that are no longer needed can be deleted.

### *Deleting Old Archive Log Files*

Archive log files will continue to be accumulated until space runs out. You would only wish to retain any archive logs that could be successfully applied to one of your backups. The following command would delete all archive log files that are older than two weeks from the current date.

```
RMAN> DELETE ARCHIVELOG ALL COMPLETED BEFORE 'SYSDATE-14';
```

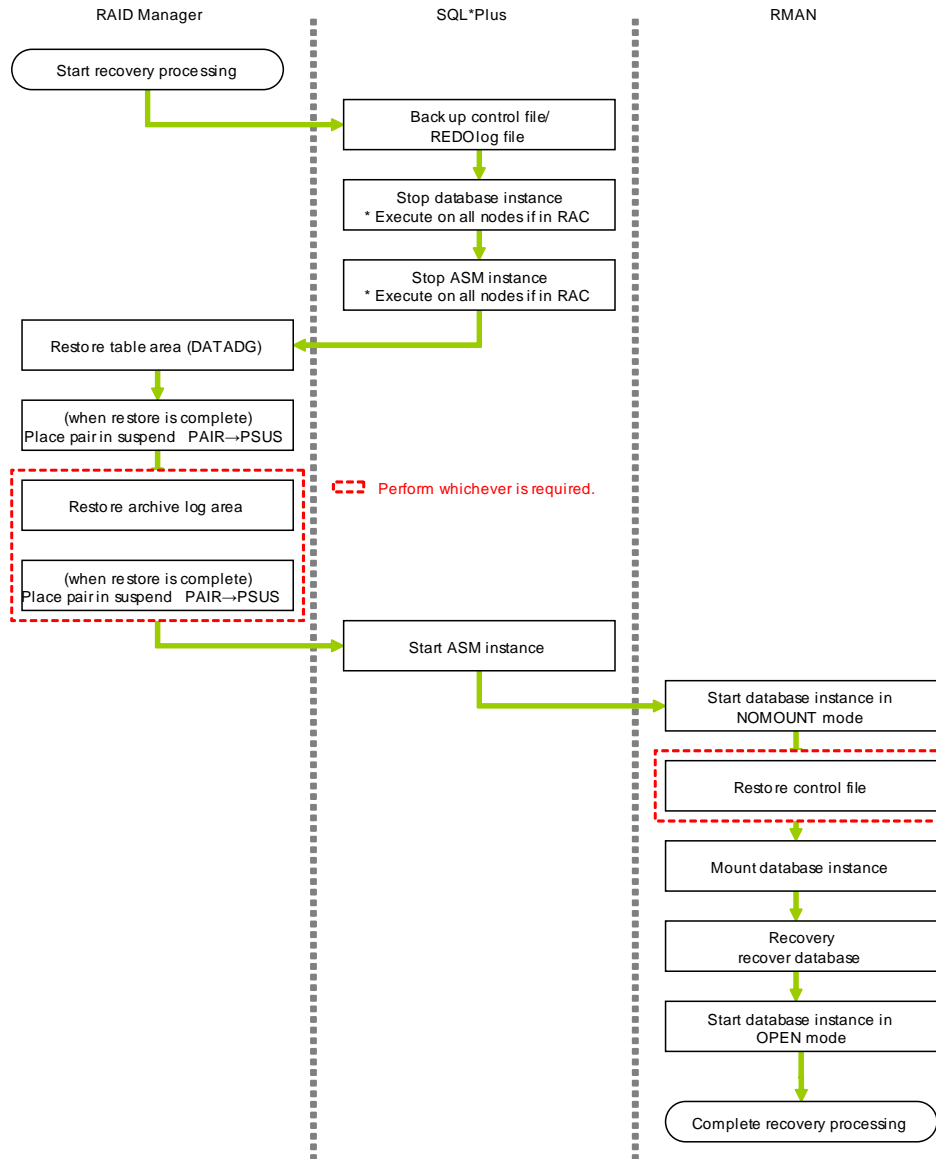
## Recovery

ShadowImage enables high-speed recovery and restoration from a hot backup copy of Oracle Database. Downtime is minimized as changed blocks are managed by differential bitmaps and only the minimal data is restored in the P-VOL.

Careful attention is required when restoring the REDO log files (REDODG) and archive logs (ARCHDG) that are left on the business server. REDODG and ARCHDG contain transactions committed right up until a fault occurs, so a full recovery may not be possible. If REDODG is damaged, you have to recreate the disk group and open the database using RESETLOGS. If ARCHDG is damaged, you can restore the disk group that has been replicated in ShadowImage.

Also, ASM disk groups might be damaged if database files are restored while an ASM instance is running.

Figure 3. Flowchart of Recovery Procedures



### Scenario 1: Hot Backup of Oracle 11g Database with ASM Using Hitachi ShadowImage Software and Oracle RMAN

This section provides guidelines for the procedures involved in the hot backup of Oracle 11g Database with ASM using Hitachi ShadowImage and Oracle RMAN. The procedures are:

- Changing an Oracle 11g Database with ASM that is running on a business server into hot backup mode
- Using ShadowImage to create a secondary volume (S-VOL) with an at-time split option to preserve consistency between volumes
- Using Oracle RMAN on the backup server to back up from the S-VOL to the file system

A copy that has been made using ShadowImage can be backed up without affecting the business server or the database.

The example below uses the smallest configuration of a three disk group environment (DATADG, ARCHDG, REDODG). It should be noted that ShadowImage always uses the suspend method.

### *Create and Split ShadowImage Pairs*

Create a ShadowImage consistency group pair on the business server and then suspend it (pairsplit).

```
# paircreate -g DATADG -vl -m grp
# paircreate -g ARCHDG -vl -m grp

# pairsplit -g DATADG
# pairsplit -g ARCHDG
```

On the backup server, edit the initialization parameter file so that the database can be mounted using the startup control file (control\_start) that was created during the backup procedure.

```
$ vi $ORACLE_HOME/initdb11asm.ora
control_files = '+ARCHDG/control_file/control_start'
```

### *Create a Consistent Image of the Database*

The following operation is done on the business server.

Verify that ASM is not rebalancing by running the following command. If rebalancing operation is not taking place, a row will not be selected.

```
$ ORACLE_SID=+ASM1 sqlplus /as sysdba
SQL> select operation, state, group_number, power from v$asm_operation;
record was not selected.
```

Resynchronize (pairresync) the DATADG and ARCHDG pair.

```
# pairresync -g DATADG
# pairresync -g ARCHDG
```

Check that the pair is in pair status (PAIR).

```
# pairevtwait -g DATADG -s pair -t 3600
# pairevtwait -g ARCHDG -s pair -t 3600
```

Perform a log switch, ensuring that it does not take place during a backup.

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

Place the Oracle on the business server in backup mode (begin backup).

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

Split DATADG into a pair (pairsplit). As the pair was created using a consistency group, it will preserve consistency after the split.

```
# pairsplit -g DATADG
```

Check that it is in suspend status.

```
# pairevtwait -g DATADG -s psus -t 3600
```

End the Oracle backup mode (end backup).

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

Perform the log switch.

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

Create two backups of the control file. Use one (control\_start) to place the Oracle into MOUNT status on the backup server. Use the other one as the control file that will be backed up.

```
RMAN> run {  
    allocate channel ctl_file type disk;  
    copy current controlfile to'+ARCHDGC/control_file/control_start';  
    copy current controlfile to'+ARCHDGC/control_file/control_backup';  
    release channel ctl_file;  
}
```

Store the backup information in the recovery catalog. This allows you to store the latest archive log in RMAN.

```
RMAN> resync catalog;
```

Split the ARCHDGC with the latest archive log into a pair.

```
# pairsplit -g ARCHDGC
```

Check that it is in suspend status.

```
# pairevtwait -g ARCHDGC -s psus -t 3600
```

### *Delete Older Archive Log Files*

Delete the archive log files that were created at least two weeks prior to the current date and are no longer needed.

```
RMAN> delete archivelog all completed before 'SYSDATE-14';
```

That completes the processing on the business server.

### *Backup from Secondary Volumes to a File System or Other Media*

This section describes processing on the backup server. We will back up from the ShadowImage S-VOL that was created on the business server to a file system.

When the S-VOL is first recognized by the backup server, the raw device name and Hitachi Dynamic Link Manager software name are sometimes different from the names in the business server. Even when this happens, the ASM instance can find the correct ASM disk group.

Start the ASM instance.

```
$ export ORACLE_SID=+ASM  
$ sqlplus /as sysdba  
SQL> startup
```

Use the S-VOL on the backup server to mount the database instance. Use the file that you created in your prior preparations (\$ORACLE\_HOME/dbs/initdb11asm.ora) as the initialization parameter file. As specified in this initialization parameter file, use the +ARCHDGC/control\_file/control\_start control file to mount the initialization parameter file.

```
$ export ORACLE_SID=db11asm  
$ sqlplus /as sysdba
```

```
SQL> startup mount pfile='$ORACLE_HOME/dbs/initdb11asm.ora'
```

Use RMAN to back up the table areas, archive REDO logs and control file to a file system.

```
RMAN> run {
  allocate channel t1 device type disk format '/tmp/%U';
  backup as copy database plus archivelog;
  backup as copy controlfilecopy '+ARCHDG/control_file/control_backup';
  release channel t1;
}
```

Finally, stop the database instance and the ASM instance.

```
SQL> shutdown
```

```
$ export ORACLE_SID=+ASM
```

```
$ sqlplus /as sysdba
```

```
SQL> shutdown
```

This completes the processing on the backup server.

When required, you can copy the files that have been backed up to the file system to another medium.

## Scenario 1b: Cold Backup of Oracle 11g Database with ASM Using Hitachi ShadowImage Software and Oracle RMAN

This section describes the procedures for creating a copy of a business server database after it has been shut down. The copy can be performed quickly, because the volumes are being copied using ShadowImage. There is no need for a consistency group (-m grp split option), but it is still recommended as a best practice. The ShadowImage pair group can include not just data files, but also control files and online REDO log files that could not be included in a hot backup process.

### *Create and Split ShadowImage Pairs*

Create a ShadowImage consistency group pair on the business server. After creating the pair, suspend it (pairsplit).

```
# paircreate -g SIALLDG -vl
```

```
# pairsplit -g SIALLDG
```

### *Create a Consistent Image of the Database*

Resynchronize the pair (pairresync) if the pair status is suspend (PSUS).

```
# pairresync -g SIALLDG
```

Stop the database instance first and then stop the ASM instance.

```
$ export ORACLE_SID=db11asm
```

```
$ sqlplus / as sysdba
```

```
SQL> shutdown
```

```
$ export ORACLE_SID=+ASM
```

```
$ sqlplus / as sysdba
```

```
SQL> shutdown
```

Perform a pair split (pairsplit) and place it in suspend status (PSUS).

```
# pairsplit -g SIALLDG
```

Verify that it is in suspend status (PSUS).

```
# pairevwait -g SIALLDG -s psus -t 3600
```

This turns the ShadowImage secondary volume into a valid copy of the database. The copy that you have created can be used for high-speed recovery to a business server, for creating a backup using RMAN on a backup server, or to create a test environment on another server.

## Scenario 2: Recovery of Oracle 11g Database with ASM Using Hitachi ShadowImage Software and Oracle RMAN

This section describes the steps involved in the restoration of corrupt or missing data files from an earlier Hitachi ShadowImage hot backup process. After the restoration of data files, the database is recovered by using Oracle Recovery Manager (RMAN) to roll-forward logs to a desired point in time.

### *Prepare Database for Recovery*

Back up the REDO log file and control file that contain the latest information.

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

```
RMAN> backup as copy current controlfile;
```

Stop the database instance and the ASM instance in order to restore the disk group in ShadowImage. In the event of RAC, stop the database instances on all nodes.

```
$ export ORACLE_SID=db11asm
```

```
$ sqlplus /as sysdba
```

```
SQL> shutdown
```

```
$ export ORACLE_SID=+ASM
```

```
$ sqlplus /as sysdba
```

```
SQL> shutdown
```

The ASM instance cannot be shut down if other database instances that use this ASM instance are still running. In that case, you can dismount the disk group for the individual database instance that you want to restore.

```
SQL> alter diskgroup DATADG dismount;
```

### *Restore Database from the S-VOLs*

Perform a pair restore (pairresync -restore) of the table area.

```
# pairresync -g DATADG -restore
```

Check that the pair is in pair status (PAIR).

```
# pairevwait -g DATADG -s pair -t 36000
```

Place the pair in suspend status (PSUS), as the contents of the secondary volume may be changed if they are still in pair status (PAIR).

```
# pairsplit -g DATADG
```

Verify that it is in suspend status (PSUS).

```
# pairevwait -g DATADG -s psus -t 3600
```

Start the ASM instance. If the disk group is not mounted, mount it.

```
$ export ORACLE_SID=+ASM
$ sqlplus /as sysdba
SQL> startup
```

or

```
SQL> alter diskgroup DATADG mount;
```

Start the database instance with NOMOUNT.

```
RMAN> startup nomount;
```

If the control file of the database instance is broken, restore it now.

```
RMAN> restore controlfile from '/tmp/cf_D-D8SI_id-2510123866_9gjjah05';
```

Start the restored database on the business server using MOUNT.

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

Perform a complete recovery or a point-in-time recovery. The following command is used for complete recovery.

```
RMAN> recover database;
```

Alternatively, you can specify a point-in-time recovery using "until SCN" or "until time".

```
RMAN> run {
    set until time '12-may-06 15:30';
    recover database;
}
```

Start the database.

```
$ RMAN> alter database open;
```

If you have recreated an incomplete recovery and control file, start the database using RESETLOGS. After you have performed RESETLOGS, obtain the backup and then resume business.

```
$ RMAN> alter database open resetlogs;
```

NOTE: RESETLOGS starts a new sequence of log numbers. In earlier versions of Oracle this would result in new logs overwriting previous archive logs limiting you from further restore/recover attempts. If you follow Oracle naming guidelines in the current version this will not happen as the log sequences will have different names. During recovery the Oracle database will know the appropriate log sequences to apply. However, you will not be able to apply logs from a different sequence, which is needed only in rare circumstances. For example, if a database gets corrupted, but you are not sure when it occurred, you won't know the exact backup that is to be restored. Such a scenario requires performing multiple restore and recovery sequences. The simplest strategy is to restore and recover backwards in time sequence, using the newest backup first, not from the oldest backup forward.

This completes recovery processing. You can start other nodes when necessary in the case of Oracle RAC. You may need to add a TEMP tablespace in some situations, such as if you wish to perform queries on the backup server.

## Scenario 3: Cloning of Oracle 11g Database with ASM Using Hitachi ShadowImage Software and Oracle RMAN

This section discusses the procedure for creating a clone of an Oracle 11g Database with ASM without using hot backup mode. It is also called a point-in-time copy. You can use this cloned copy to create an exact copy of the production database for use in development or testing environments. A data file, control file, REDO log file and Flashback log file must be included in the ShadowImage group in order to create a clone. The cloned database is created using the ShadowImage consistent at-time split function.

### *Create ShadowImage Pair*

Create a pair(paircreate) using the "-m grp" option to activate the at-time split function.

```
# paircreate -g SIALLDG -vl -m grp
```

When the pair has been created, split it (pairsplit).

### *Clone the Database*

Resynchronize the pair (pairresync).

```
# pairresync -g SIALLDG
```

Verify that the pair is in pair status (PAIR).

```
# pairevwait -g SIALLDG -s pair -t 3600
```

Split the pair (pairsplit) while the database instance is running. The consistency of the pair group will be preserved during the split because it was created in a consistency group.

```
# pairsplit -g SIALLDG
```

Verify that the pair is in suspend status (PSUS).

```
# pairevwait -g SIALLDG -s psus -t 3600
```

Start an ASM instance on the test server. Check that the correct permissions have been set for the ShadowImage secondary volume on the test server, and that the specifications for ASM\_DISKSTRING and ASM\_DISKGROUPS are correct.

```
$ export ORACLE_SID=+ASM
$ sqlplus /as sysdba
SQL> startup
```


Mount the cloned database on the test server.

```
$ export ORACLE_SID=db11asm
$ sqlplus /as sysdba
SQL> startup mount
SQL> recover database;
SQL> exit
```

Rename the cloned database (that is, change ORACLE\_SID) by using the nid utility.

```
$ nid target=sys/manager1@test DBNAME=oratest
```

Open the cloned database.



```
$ export ORACLE_SID=oratest  
$ sqlplus /as sysdba  
SQL> startup mount  
SQL> alter database open;
```

You can now use the database that has been replicated on the test server. You need to recover the database, as its data is in the same state as if the database had been forcibly shut down by a power failure.



## Conclusion

The objective of this document was to provide guidelines for optimizing the backup and recovery process of an Oracle 11g Database with ASM using Hitachi storage solutions. The combination of the Hitachi Universal Storage Platform V or Hitachi Adaptable Modular Storage 2000 family system with Hitachi ShadowImage Heterogeneous Replication or ShadowImage Replication software and Oracle RMAN provides an optimized backup and recovery system for Oracle 11g Database with ASM. Enterprises can deploy a cost-effective backup and recovery solution for Oracle 11g Database with ASM by using the guidelines provided in this document.

## Schedule A: RMAN Recovery Catalog Guidelines

This page presents the commands frequently used by Oracle Recovery Manager (RMAN) and examples of their use.

The recovery catalog can manage the database on the business server and the database on the backup server. The methods for doing this are shown below, as you need to connect to the target recovery catalog, irrespective of whether it is on the business server or the backup server.

Connecting to the database and recovery catalog (rcat) on the business server (prod):

```
$ rman target system/manager1@prod catalog rman/rman@rcat
```

Connecting to the database and recovery catalog (rcat) on the backup server (prod):

```
$ rman target system/manager1@bkup catalog rman/rman@rcat
```

An example of creating the recovery catalog is shown below.

```
$ rman catalog rman/rman@rcat
```

```
RMAN> create catalog;
```

```
RMAN> exit
```

Connect to the database and the recovery catalog.

```
$ rman target prod catalog rman/rman@rcat
```

Register the database on the connected recovery catalog.

```
RMAN> register database;
```

**Database registered in recovery catalog**

**Starting full resync of recovery catalog**

**Full resync complete**

Check that it is registered properly.

```
RMAN> report schema;
```

```
List of Permanent Datafiles
=====
File Size(MB) Tablespace RB segs Datafile Name
-----
1 490 SYSTEM YES +DATADG/db11asm/datafile/system.256.617205169
2 25 UNDOTBS1 YES +DATADG/db11asm/datafile/undotbs1.258.617205171
3 310 SYSAUX NO +DATADG/db11asm/datafile/sysaux.257.617205169
4 5 USERS NO +DATADG/db11asm/datafile/users.259.617205171
5 100 EXAMPLE NO +DATADG/db11asm/datafile/example.261.617205255
List of Temporary Files
=====
File Size(MB) Tablespace Maxsize(MB) Tempfile Name
-----
1 20 TEMP 32767 +DATADG/db11asm/tempfile/temp.260.617205251
```

## Schedule B: Example of the TNSNAMES.ORA Configuration File

An example of the tnsnames.ora file is shown below.

### **tnsnames.ora**

```
#-----  
PROD=  
  (DESCRIPTION=  
    (ADDRESS_LIST=  
      (ADDRESS= (PROTOCOL=TCP) (PORT=1521) (HOST=PRODHOST) ) )  
    (CONNECT_DATA=  
      (SID_NAME=PROD) )  
  )  
BKUP=  
  (DESCRIPTION=  
    (ADDRESS_LIST=  
      (ADDRESS= (PROTOCOL=TCP) (PORT=1521) (HOST=BKUPHOST) ) )  
    (CONNECT_DATA=  
      (SID_NAME=BKUP) )  
  )  
RCAT=  
  (DESCRIPTION=  
    (ADDRESS_LIST=  
      (ADDRESS= (PROTOCOL=TCP) (PORT=1521) (HOST=BKUPHOST) ) )  
    (CONNECT_DATA=  
      (SID_NAME=RCAT) )  
  )
```

## Schedule C: Example of Hitachi ShadowImage Configuration Files

Examples of ShadowImage configuration files are shown below. These support the primary and secondary volumes.

### **/etc/horcm0.conf**

```
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
db11asm1 horcm0 12000 3000

HORCM_CMD
#dev_name dev_name dev_name
/dev/rdisk/c3t50060E80102002F4d0s2

HORCM_DEV
#dev_group dev_name port# TargetID LU# MU#
DATADG ora_dev00 CL1-A-1 0 22 0
DATADG ora_dev01 CL1-A-1 0 23 0
DATADG ora_dev02 CL1-A-1 0 24 0
ARCHDG ora_dev03 CL1-A-1 0 25 0
ARCHDG ora_dev04 CL1-A-1 0 26 0
ARCHDG ora_dev05 CL1-A-1 0 27 0
REDODG ora_dev06 CL1-A-1 0 28 0
SIALLDG ora_dev10 CL1-A-1 0 22 0
SIALLDG ora_dev11 CL1-A-1 0 23 0
SIALLDG ora_dev12 CL1-A-1 0 24 0
SIALLDG ora_dev13 CL1-A-1 0 25 0
SIALLDG ora_dev14 CL1-A-1 0 26 0
SIALLDG ora_dev15 CL1-A-1 0 27 0
SIALLDG ora_dev16 CL1-A-1 0 28 0

HORCM_INST
#dev_group ip_address service
DATADG db11asm1 horcm1
ARCHDG db11asm1 horcm1
REDODG db11asm1 horcm1
SIALLDG db11asm1 horcm1
```

### **/etc/horcm1.conf**

```
HORCM_MON
#ip_address service poll(10ms) timeout(10ms)
db11asm1 horcm1 12000 3000

HORCM_CMD
#dev_name dev_name dev_name
/dev/rdisk/c3t50060E80102002F4d0s2

HORCM_DEV
#dev_group dev_name port# TargetID LU# MU#
DATADG ora_dev00 CL1-B-4 0 29 0
DATADG ora_dev01 CL1-B-4 0 30 0
DATADG ora_dev02 CL1-B-4 0 31 0
ARCHDG ora_dev03 CL1-B-4 0 32 0
ARCHDG ora_dev04 CL1-B-4 0 33 0
ARCHDG ora_dev05 CL1-B-4 0 34 0
REDODG ora_dev06 CL1-B-4 0 35 0
SIALLDG ora_dev10 CL1-B-4 0 29 0
SIALLDG ora_dev11 CL1-B-4 0 30 0
SIALLDG ora_dev12 CL1-B-4 0 31 0
SIALLDG ora_dev13 CL1-B-4 0 32 0
SIALLDG ora_dev14 CL1-B-4 0 33 0
SIALLDG ora_dev15 CL1-B-4 0 34 0
SIALLDG ora_dev16 CL1-B-4 0 35 0

HORCM_INST
#dev_group ip_address service
DATADG db11asm1 horcm0
ARCHDG db11asm1 horcm0
REDODG db11asm1 horcm0
SIALLDG db11asm1 horcm0
```

## Schedule D: Example of an Initialization Parameter File

The key parameters in the initialization parameter file for a database instance and ASM instance are shown below.

### *Business Server*

```
init+ASM.ora
INSTANCE_TYPE=ASM
ASM_DISKSTRING='/dev/sddl*'  
ASM_DISKGROUPS='DATADG, ARCHDG, REDODG'
```

```
initdb11asm.ora  
*.db_name='db11asm'  
*.control_files=('+FALSHDG/db11asm/control01.ctl', '+REDODG/db11asm/control02.ctl')  
*.log_archive_dest_1='LOCATION=+ARCHDG/'
```

### *Backup Server*

```
init+ASM.ora
INSTANCE_TYPE=ASM
ASM_DISKSTRING='/dev/sddl*'  
ASM_DISKGROUPS='DATADG, ARCHDG'
```

```
initdb11asm.ora  
*.db_name='db11asm'  
*.control_files='+ARCHDG/control_file/cold_control_start'  
*.log_archive_dest_1='LOCATION=+ARCHDG/'
```



---

**Corporate Headquarters** 750 Central Expressway, Santa Clara, California 95050-2627 USA  
Contact Information: + 1 408 970 1000 [www.hds.com](http://www.hds.com) / [info@hds.com](mailto:info@hds.com)

**Asia Pacific and Americas** 750 Central Expressway, Santa Clara, California 95050-2627 USA  
Contact Information: + 1 408 970 1000 [www.hds.com](http://www.hds.com) / [info@hds.com](mailto:info@hds.com)

**Europe Headquarters** Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom  
Contact Information: + 44 (0) 1753 618000 [www.hds.com](http://www.hds.com) / [info.emea@hds.com](mailto:info.emea@hds.com)

Hitachi is a registered trademark of Hitachi, Ltd., and/or its affiliates in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

ShadowImage and TrueCopy are registered trademarks and Universal Storage Platform is a trademark of Hitachi Data Systems Corporation.

All other trademarks, service marks and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect and that may be configuration dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to <http://www.hds.com/corporate/legal/index.html> or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

© Hitachi Data Systems Corporation 2008. All Rights Reserved.

WHP-300-00 DG October 2008