



Oracle Database 10g Automatic Storage Management Best Practices with Hitachi Replication Software on the Hitachi Universal Storage Platform™ Family of Products

Application Brief

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

June 2007

Executive Summary

Today's enterprises search for backup solutions that are high in quality *and* cost-effective. In answer to these requirements, Hitachi Data Systems and Hitachi, Ltd., have tested and evaluated best practices for using Automatic Storage Management (ASM), a new feature introduced in Oracle Database 10g, to simplify Oracle database file administration on the Hitachi Universal Storage Platform™ family of products while optimizing performance.

Oracle Database 10g ASM provides storage cluster volume management and file system functionality at no additional cost. ASM increases storage utilization, performance, and availability while eliminating the need for third-party volume management and file systems for Oracle database files. As a result, ASM provides significant cost savings for the data center.

In addition to pure ASM implementations, ASM co-exists and operates in mixed third-party volume manager and file system environments that are needed to store non-Oracle database files.

ASM-demonstrated superior capabilities in terms of database file automation, reduction of complexity, and improved performance mesh perfectly with the industry-leading Hitachi Universal Storage Platform (USP). The USP offers great value for storage consolidation, virtualization, and simplification of storage management. Together, ASM and the USP provide robust database reliability, rock-solid backup and disaster recovery, and robust management options.

Contents

Introduction	1
Hitachi Universal Storage Platform V Overview	3
Universal Storage Platform Architecture.....	3
Control Memory/Shared Memory Package.....	3
Cache Memory Packages	4
Channel Adapter for Front-end Director.....	4
Disk Adapter for Back-end Director	4
Disk Drives	4
Oracle Database 10g—ASM Feature Overview	5
ASM Striping	5
ASM Rebalancing	5
ASM Redundancy.....	6
Using ASM with Universal Storage Platform Volumes	6
Universal Storage Platform Recommended Settings	6
Hitachi Dynamic Link Manager Software Overview.....	6
Hitachi ShadowImage Heterogeneous Replication Software Overview	7
Hitachi TrueCopy Asynchronous and Hitachi Universal Replicator Software Overview.....	7
Consistency Groups and At-Time Split Options for Point-in-Time Copies	10
ASM Recommended Settings	11
Creating an ASM Instance	11
Creating ASM Disk Groups.....	12
Configuring ASM Disk Groups and CRS Volumes	12
Using ASM Instances.....	14
Using ASM Disk Groups	16
ASM Files	17
ASM File Information Acquisition Method	17
ShadowImage Software Recommended Settings.....	18
Pair Configuration	18
Pair Control	18
Command Descriptions	20

ASM Backup Using Hitachi ShadowImage, TrueCopy, or Universal Replicator Software.....	20
Backup Overview.....	20
Comparing Oracle Database 10g and Oracle9i Database	21
Cold Backup Precautions	22
Hot Backup Precautions.....	22
Backup Requirements for Specific Data Types.....	22
Recovery Overview	23
Scenario 1: Backup of Oracle Database 10g ASM Using ShadowImage Software	25
Scenario 2: Recovery of Oracle Database 10g ASM Using ShadowImage Software	28
Scenario 3: Oracle Database 10g ASM Cloning Using ShadowImage Software	29
Scenario 4: Oracle Database 10g Cold Backup Using ShadowImage Software's Overall Flowchart of Backup Procedures.....	30
Appendix A: RMAN Recovery Catalog Guidelines	33
Appendix B: Sample TNS Configuration File	35
Appendix C: Sample ShadowImage HORCM Configuration Files	36
Appendix D: Sample Contents of ASM Instance Parameter File.....	38

Oracle Database 10g Automatic Storage Management Best Practices with Hitachi Replication Software on the Hitachi Universal Storage Platform™ Family of Products

Application Brief

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

Introduction

This application brief describes guidelines and best practices for Oracle Database 10g backup and recovery using the Automatic Storage Management (ASM) feature with Hitachi replication technologies. The ASM feature on the Oracle Database 10g simplifies storage and administration tasks for enterprise IT organizations.

Hitachi Data Systems and Hitachi, Ltd., have tested and evaluated best practices for using ASM on the Hitachi Universal Storage Platform™ as well as the latest version, the Hitachi Universal Storage Platform™ V, (also referred to as the “USP” and the “USP V”, respectively, in this document). This best practices application brief explores testing and practices that can support cost-effective, high-quality backup, cloning, and recovery solutions based upon Oracle Database 10g with ASM. Hitachi Data Systems provides for in-system point-in-time (PiT) replication using Hitachi ShadowImage™ Heterogeneous Replication or Hitachi Copy-on-Write software, and remote storage system replication using Hitachi TrueCopy® Synchronous, Hitachi TrueCopy Asynchronous, or Hitachi Universal Replicator software. In addition to providing protection for a central site, these replication products can be combined to provide extra levels of business continuity options at local and remote disaster recovery sites.

Although the functional procedures for backup/recovery of Oracle databases are similar whether using ASM or regular file systems, there are distinct procedures and configurations required by the specific software program as described in this application brief.

PiT copies using Hitachi ShadowImage, Copy-on-Write, TrueCopy, or Universal Replicator software allow for copies of data to be created quickly. This application brief will focus on showing how to create ShadowImage and TrueCopy copies of an Oracle Database 10g ASM for:

- Database backup
- Database recovery
- Cloning databases for test and development

Hardware replication within the Hitachi storage platform is more efficient than server-based replication because it offloads the processing from the server to the storage controller. In addition, centralized management of replication at the storage controller can help keep administrative costs under control.

Oracle Database 10g provides a built-in recovery solution called Flashback Database for a continuous PiT copy solution. Oracle Flashback Database allows you to quickly recover an Oracle database to a given, specific time to correct problems caused by logical data corruptions or user errors. By rolling back to a specific time, the database can be restored to an error-free state. Then, updates can be selectively applied to bring the database current.

Both Hitachi storage replication and Oracle Flashback can be utilized to offer independent and multiple levels of protection and recovery options.

Hitachi replication products can be used to meet the requirements of Oracle Database 10g ASM split-mirror guidelines, which include:

- Creating PiT split copies of the required ASM disk group LUNs
- Enabling all LUNs in a disk group to be split atomically with respect to write completion, thus preserving write order fidelity by using Hitachi consistency group technology
- Creating two ASM disk groups that are replicated along with a third disk group for redo logs that is not required to be replicated
 - “DATADG” disk group contains data files only (no redo logs or control files)
 - “FLASHDG” disk group contains archive logs, flashback logs, control files
 - “REDODG” disk group contains online redo logs
- The Hitachi replication groups contain all the volumes in the Oracle Database 10g ASM disk groups
- The backup host is configured with the proper operating system (OS) environment, Oracle binaries, init.ora/password files, directories, and userids/groups similar to the production host
- The backup host has Oracle Net connectivity to the Oracle Recovery Manager (“RMAN”) recovery catalog

While this application brief for convenience refers primarily to the latest and best-performing Hitachi flagship product, the USP V, due to architectural, microcode, and software similarities, these best practices apply equally to the entire product line unless otherwise noted:

- Hitachi Universal Storage Platform V
- Hitachi Universal Storage Platform
- SANRISE Universal Storage Platform (largely sold and supported in the Japan market by Hitachi, Ltd., rather than Hitachi Data Systems)

- Hitachi Network Storage Controller™ (NSC)
- Hitachi Adaptable Modular Storage systems

Oracle ASM provides superior capabilities in terms of database file automation, reduction of complexity, and improved performance as compared to third-party volume managers. The industry-leading Hitachi Universal Storage Platform offers great value for storage consolidation, virtualization, and simplification of storage management.

By implementing both industry-leading platforms together, the enterprise can truly realize the best of all worlds. The USP's native abilities to optimize performance and to seamlessly replicate/backup data in high-volume production environments perfectly complements ASM. Together, ASM and the USP provide robust database reliability, rock-solid backup and disaster recovery, and robust management options.

This application brief both presents and reconciles Best Practices for each platform. In some cases, each platform's Best Practices, taken standalone, need to be modified or overridden when both products are implemented together. These cases are noted and described in this document.

Hitachi Universal Storage Platform V Overview

The Hitachi Universal Storage Platform V is powered by an enhanced Hitachi Universal Star Network™ V crossbar switch architecture with advanced processor algorithms and new workload distribution amongst multiprocessors, 4Gb/sec host ports, 4Gb/sec disks, and 4Gb/sec switched back-end directors. It provides 40 percent more peak performance than and further extends the scalability and enterprise-class functionality of the previous USP version. The latest-generation Universal Storage Platform V establishes a new industry category with previously unattainable levels of consolidation and virtualization of up to 332TB internal and 247PB heterogeneous external storage in one pool and supports application-centric storage management and simplified data replication. The Universal Storage Platform V enables storage managers to logically partition storage resources to maximize application quality of service (QoS), move data across tiers of storage to match application attributes, establish enterprise-wide business continuity, and manage it all from a “single pane of glass.”

Universal Storage Platform Architecture

The highly reliable Hitachi Universal Storage Platform architecture uses components (front-end directors, back-end directors, cache, cache switch, and internal paths) that are fully multiplexed and redundant. Additionally, in the latest USP V, the front-end director, cache, and back-end director paths have 40 percent more system performance than the previous USP. The maximum internal data transmission capacity of cache memory has been increased more than six-fold to a maximum of 106GB/sec. The unique and independent control memory network bandwidth, for the exchange of control memory and control information, has been increased significantly. This facilitates the high performance you need to handle a large-scale storage consolidation environment.

Control Memory/Shared Memory Package

Control or shared memory is memory that can be accessed by each front-end director and back-end director. It stores the control information (cache memory control information and the like) that is necessary to access user data. In previous systems, shared memory was installed on top of the cache memory package, but in the Universal Storage Platform product family, shared memory is an independent hardware package.

Packages housing shared memory are typically installed package by package with front and back sides forming clusters. By so doing, in the unlikely event that there is a failure in one side of a package, the information stored on the other side of the package can still be accessed. Additionally, shared memory packages can be installed to a maximum of four packages (two front sides and two back sides). In previous devices, when a write-through operation was performed in the event of a cache memory package failure, the system would wait for write processing to the disk drive to be completed and then report to the host that writing has been completed. From a data availability standpoint, when the control information configuration is installed in a four-sided package, an operation is possible in which a completion report is made to the host with a normal cache memory write operation—without performing a write-through operation. By doubling the data paths to shared memory, high-speed access to the control information needed for access to user data and operations is enabled, and the performance of all storage systems can be enhanced. At the same time, the effects of package failures can be minimized.

Cache Memory Packages

Cache memory is memory used to perform data read/write processing efficiently between front-end and back-end directors. The cache capacity on both sides is configured to be a minimum of 4GB to a maximum of 256GB, with expansion in units of 4GB or 8GB. The disk controller performs control so that there is always the most efficient use of the cache in response to data access patterns, thus obtaining a highly stable level of performance. As an optional expansion, access paths to the cache can be doubled to obtain internal data transmission capacities in combination with a cache switch for a maximum of 106GB/sec.

Channel Adapter for Front-end Director

This is the control path that controls data transmission between the host interface (channel) and cache memory. Although the number of control processors installed in the front-end director depends on the types of interfaces, a maximum of eight control processors are installed in a single front-end director package. Front-end directors are expanded in units of two array control frames, and a maximum of 12 units can be installed, although this depends on the number of back-end directors and cache switches installed. The maximum number of connection ports supported on the USP V is 224 Fibre Channel ports or 112 IBM® FICON® ports.

Disk Adapter for Back-end Director

The disk adapter package controls the transfer of data between cache memory disk services. A 400MB/sec Fibre Channel Arbitrated Loop (FC-AL) interface is used as the disk drive interface. There are eight control processors housed in a single adapter, supporting eight drive paths. Adapters are expandable in disk controller units, and can store up to eight, depending on the number of disk drives installed. Storage systems can store a maximum total of 128 microprocessors with front-end and back-end directors taken together.

Disk Drives

The USP V supports a 400MB/sec FC-AL interface provided with dual ports, capacity of 73GB/146GB/300GB, with disk speeds of 10,025rpm (min-1) for high-speed hard disk drives and 73GB and 146GB capacity, and disk speed of 14,904rpm (min-1) for ultrahigh-speed hard disk drives. The USP basic unit (array control frame) is installed with a maximum of 128 disk drives, but an additional four array frames that can store a further 256 disk drives each can be connected. The overall system can store a maximum of 1,152 hard disk drives internally.

Oracle Database 10g—ASM Feature Overview

ASM, a new feature introduced in Oracle Database 10g, simplifies database file storage and administration tasks. ASM is a cluster volume manager and a file system for Oracle database files. Therefore, it eliminates the need for third-party volume managers and file systems for storing database files and provides significant cost savings for the IT database environments. However, ASM operates in a mixed volume management and file system environment that is necessary to store the non-Oracle database files.

In preparing this document, Hitachi, Ltd., and Hitachi Data Systems focused on the ASM's capabilities as a way for enterprises to cut build costs and to reduce build time. ASM contributes to the creation of high-quality backup solutions and offers excellent cost-effectiveness.

An ASM disk group typically consists of multiple LUNs. This allows ASM files to be evenly distributed among all the LUNs in the disk group. ASM automatically distributes the ASM file extents and rebalances I/O by redistributing blocks if LUNs are added or removed.

ASM Striping

There are two default stripe widths used with ASM. It is not traditional striping as with most file systems but rather an even I/O distribution policy within an ASM disk group. A FINE stripe width of 128K is the default for control files and redo logs. A COARSE stripe width of 1M is used for all other data files. The best practice is to leave these defaults unchanged for most database workloads.

ASM Rebalancing

When a LUN is added or removed from an ASM disk group, ASM will start rebalancing the blocks evenly among the modified ASM disk group. The database administrator can change the priority of rebalancing or turn it off if desired. This rebalancing operation can be performed concurrently with Hitachi storage-based snapshots as long as consistency groups are used for atomic splits. Using Oracle Database 10g ASM rebalancing is a very effective method of providing additional space while also increasing the dispersing of the workload evenly across LUNs. This follows Oracle's Stripe and Mirror Everything (SAME) approach to storage management. The blocks are striped by ASM while the Hitachi storage mirrors or protects with hardware RAID levels. The best practice is to evenly assign LUNs from many Hitachi RAID array groups for a given ASM disk group. Assigning all the LUNs from a single parity group to a single ASM disk group will not allow ASM rebalancing to be effective.

Oracle Hot Backup of ASM disk groups require that TrueCopy, Universal Replicator, or ShadowImage consistency groups be used for atomic splits, otherwise inconsistencies might occur due to ASM rebalance operations. As long as consistency groups are used all ASM functionality is retained, even during the Hot Backup process.

You may check to see if ASM rebalancing is occurring by issuing the following query. If any rows are returned then rebalancing is being performed.

```
$ export ORACLE_SID=+ASM
$ sqlplus / as sysdba
SQL> select operation,state,group_number,power from v$asm_operation;

no rows selected
```

Scheduling ASM Rebalancing

You may want to postpone rebalancing until periods of lower processing or after many disk group changes. You can turn on or off rebalancing and set the priority if desired. A power level of 0 turns off rebalancing. A power level of 1–11 is used to enable rebalancing with a higher value putting priority on the rebalance operation. The default value of 1 minimizes disruption to the database consuming less processing and I/O resources. As mentioned previously, rebalancing does not need to be turned off during Oracle Hot Backup using ShadowImage or TrueCopy software if consistency groups are used. For completeness, however, an example of turning off ASM disk group rebalancing follows:

```
$ export ORACLE_SID=+ASM  
$ sqlplus / as sysdba  
SQL> alter diskgroup DATADG rebalance power 0
```

ASM Redundancy

ASM disk groups support three levels of data protection. ASM disk groups can be specified with normal, high, or external redundancy. Normal provides for a two-way mirror. High provides three-way mirroring. External redundancy relies on the hardware for data protection and is the best practice on Hitachi storage. If you plan to use TrueCopy or ShadowImage software, then external redundancy is required.

Using ASM with Universal Storage Platform Volumes

Universal Storage Platform Recommended Settings

A shared volume is required when you build a Real Application Cluster (RAC) database. There are no particular settings that must be made when setting up a shared volume other than the Host Storage Domain settings for the attached OS platform such as Linux, Sun Solaris, IBM AIX, or Microsoft® Windows. It is recommended to define paths from multiple ports for the volumes that you wish to make shared. It is also recommended to define multiple ports for ASM volumes for each server and use Hitachi Dynamic Link Manager software to manage the multiple paths.

Hitachi Dynamic Link Manager Software Overview

Hitachi Dynamic Link Manager (HDLM) software manages access paths to storage systems. HDLM provides functionality for distributing the load across paths and switching to another path if there is a failure in a path being used, thus improving system availability and reliability.

HDLM has the following features:

- **Distributes the load across paths (*load balancing*)**

When multiple paths connect a host and storage system, HDLM distributes the load across the paths. This prevents a heavily loaded path from affecting processing speed.

- **Allows processing to continue if there is a failure (*failover*)**

When multiple paths connect a host and storage system, HDLM automatically switches to another path if there is a failure in a path being used. This allows processing to continue without being affected by the failure.

- **Allows you to place online a path recovered from an error (*failback*)**

When a path recovers from an error, HDLM can place the path online. This enables the maximum number of paths to be online, which in turn enables HDLM to distribute the load for each path. Route failback can be performed manually or automatically. In an automatic failback, HDLM automatically restores the route to the active state after the user has corrected the problem in the physical route.

- **Automatically checks the path status at regular intervals (*path health checking*)**

HDLM can detect errors by checking the status of the paths at user-defined time intervals. This allows you to check for any existing path errors and to resolve them accordingly.

On Linux, each path to a LUN is assigned a device name such as /dev/sda, /dev/sdb, etc. HDLM finds paths to the same actual LUN and assigns a common HDLM name such as /dev/sddlmaa. The HDLM device format on Linux is /dev/sddlm[a-p][a-p][1-15]. Partitions on an HDLM device are represented by the number after the sddlm* name. The whole device is represented without specifying the number. You must use the HDLM device name for HDLM load-balancing and path-failover functionality to be active.

For more information, refer to the appropriate HDLM Users Guide for the specific OS platform.

Hitachi ShadowImage Heterogeneous Replication Software Overview

ShadowImage software is a storage-based hardware solution that creates RAID-protected duplicate volumes within the USP or externally connected storage systems as part of the USP virtualization using Universal Volume Manager software. ShadowImage primary volumes (P-Vols) contain the original data and up to nine secondary volumes (S-Vols) can be created as copies. These ShadowImage volumes can be grouped as part of consistency groups based on application requirements for backup and recovery.

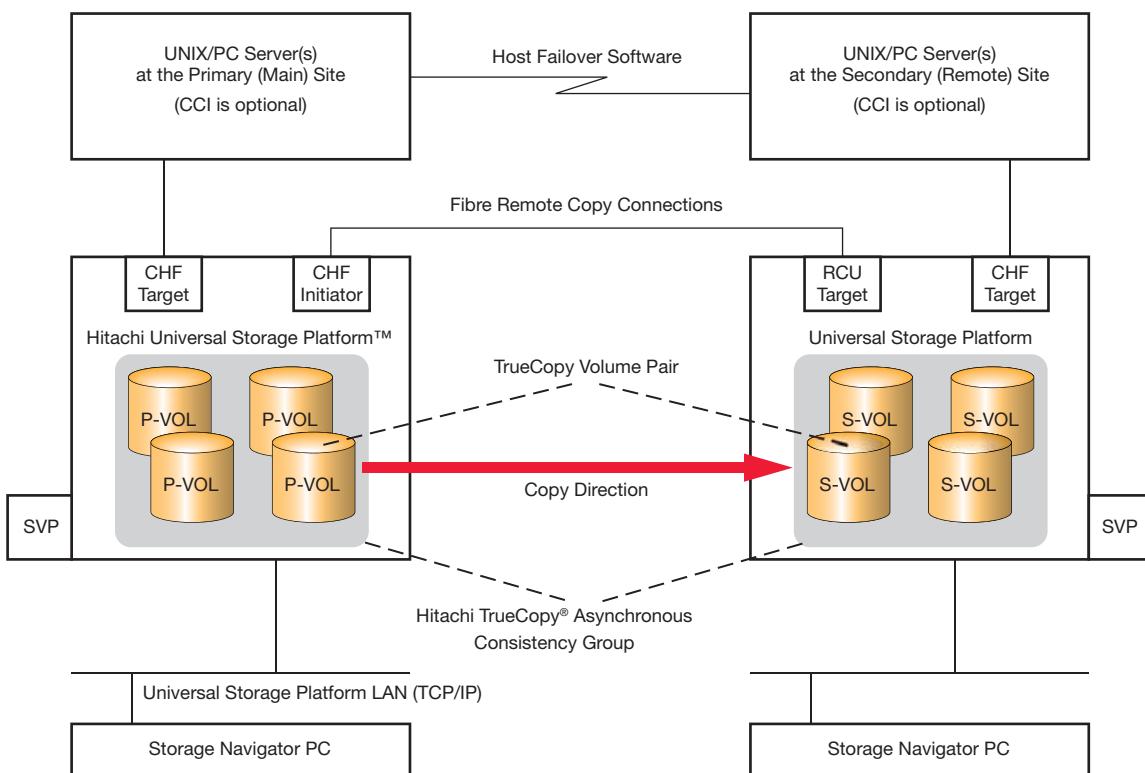
ShadowImage operations are nondisruptive and allow the primary volumes to remain online to all hosts for read and write I/O operations.

Refer to Hitachi Universal Storage Platform and Network Storage Controller ShadowImage User and Reference Guides for more information.

Hitachi TrueCopy Asynchronous and Hitachi Universal Replicator Software Overview

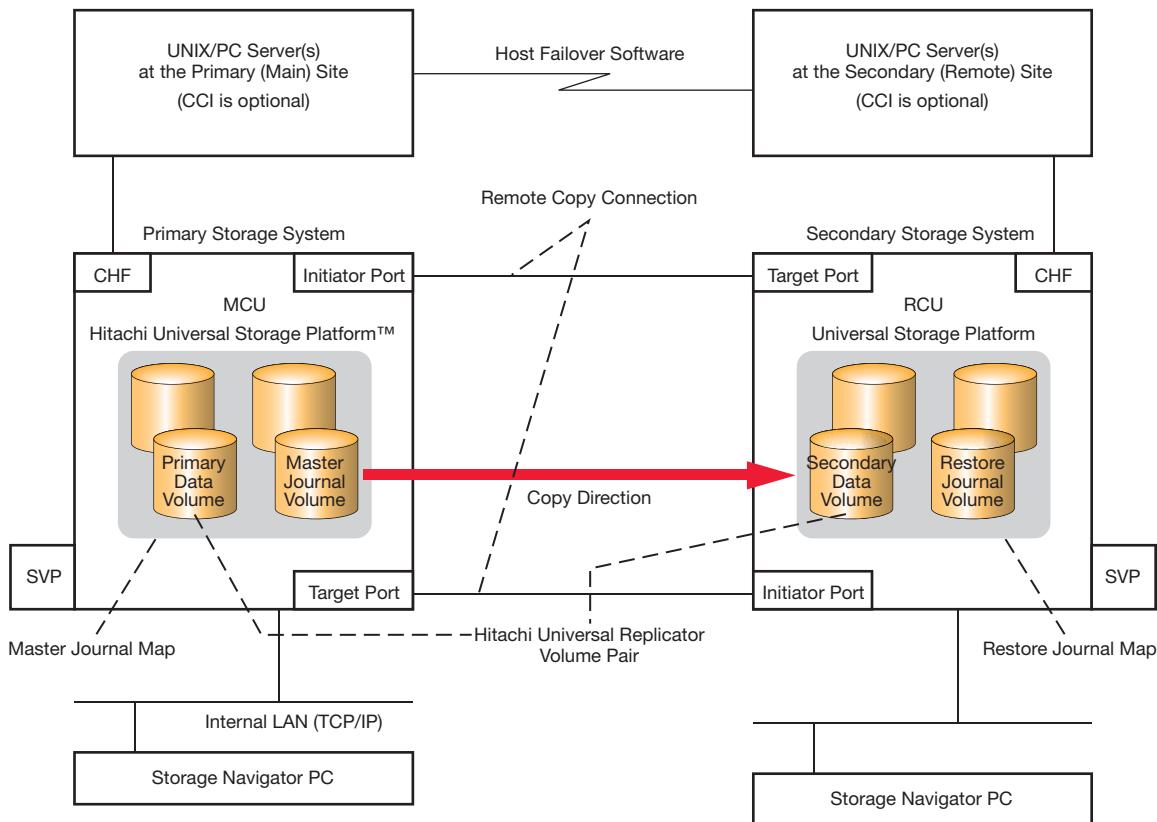
TrueCopy Synchronous and TrueCopy Asynchronous software enable remote copy operations between USP storage platforms. Remote copy operations are again nondisruptive to the primary volumes and allow hosts to continue read and write I/O operations. TrueCopy Asynchronous (TCA) software provides update sequence consistency for user-defined groups of volumes (such as large databases) as well as protection for write-dependent applications in the event of a disaster. (See Figure 1.)

Figure 1. TrueCopy Environment Example



Hitachi Universal Replicator (HUR) is an asynchronous replication technology similar to TrueCopy Asynchronous that reduces the cache requirements by utilizing disk journals to store writes during excessive peak write activity or link outages. The HUR group-based update sequence consistency solution enables fast database recovery. (See Figure 2.)

Figure 2. Example: Hitachi Universal Replicator Environment

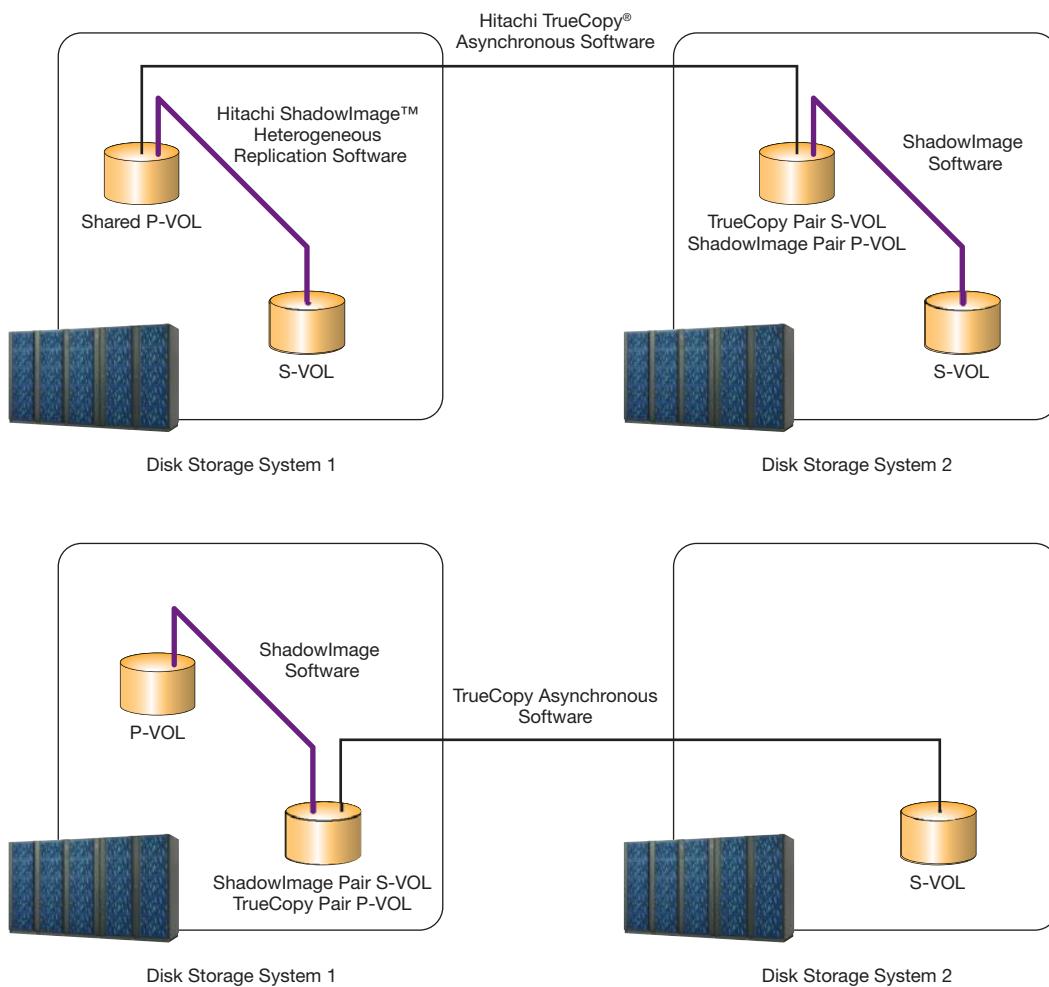


TCA and HUR represent unique and outstanding disaster recovery solutions for large amounts of data, which span multiple volumes.

Refer to the Hitachi Universal Storage Platform and Network Storage Controller TrueCopy User and Reference Guide or the Hitachi Universal Replicator User and Reference Guide for more information.

TCA or HUR operations can be combined with ShadowImage capabilities to provide multiple levels of protection and recovery options. (See Figure 3.)

Figure 3. Example: Combined ShadowImage and TrueCopy Software Environment



Consistency Groups and At-Time Split Options for Point-in-Time Copies

The ShadowImage At-Time Split function applies to ShadowImage pairs that belong to a consistency group. This function allows you to create S-VOLs of all P-VOLs in the same consistency group when the `pairsplit` command is executed using the Command Control Interface (CCI) software from the UNIX/PC server host to the Universal Storage Platform. The S-VOLs contain the same data as the P-VOLs when the Split operation is performed.

A ShadowImage consistency group is a user-defined set of ShadowImage volume pairs used for the At-Time Split function. ShadowImage consistency groups also correspond to the groups registered in the CCI configuration definition file. You may configure up to 128 consistency groups in a storage system. You may define up to 4,096 ShadowImage pairs in a consistency group. When the ShadowImage At-Time Split function is enabled, data in all P-VOLs in the same consistency group is suspended atomically in the corresponding S-VOLs at the time when the Universal Storage Platform system receives the `pairsplit` request from the host server.

The TCA and HUR group-based update sequence consistency solution enables fast and accurate database recovery, even after a “rolling” disaster, without the need for time-consuming data recovery procedures. TCA and HUR volume groups at the remote site can be recovered with full update sequence consistency, but the updates may be behind the primary site due to the asynchronous remote copy operations.

A TCA or HUR consistency group is a user-defined set of volume pairs across which update sequence consistency is maintained and ensured at the remote site. Each volume pair must be assigned to a consistency group. TCA and HUR allow you to configure up to 128 consistency groups (0-7F) for each primary storage system.

ASM Recommended Settings

In most cases, using the HDLM device name for ASM device names is preferred and specified in the `ASM_DISKSTRING` parameter and also shown in the appendices.

```
ASM_DISKSTRING='/dev/sdd1m*' 
```

This corresponds to the native Linux devices `/dev/sd*`.

ASM device names do not need to be identical on the multiple hosts in a RAC environment as long as the ASM discovery string is set appropriately and all ASM disks can be discovered. Some database administrators may prefer this practice. If so, it can be accomplished with symbolic links. In any case, identical device names are not required.

Symbolic link to an HDLM device:

```
# ln -s /dev//sdd1maa /dev/asm/lu001 
```

After creating symbolic links:

Change the ownership using the “chown” command so that the Oracle userid can access the links:

```
# chown oracle:oinstall /dev/asm/lu* 
```

Set the `ASM_DISKSTRING` to use the symbolic link prefix chosen:

```
ASM_DISKSTRING='/dev/asm/lu*' 
```

```
SQL> create diskgroup DATADG external redundancy
      DISK '/dev/asm/lu001', '/dev/asm/lu002',
      '/dev/asm/lu003' ; 
```

Creating an ASM Instance

Although ASM instances can be created manually, we recommend that they be created by Database Configuration Assistant (DBCA). Do not create a new ASM instance if one has already been created.

Creating ASM Disk Groups

You need to create an ASM disk group (DG) in order to build an Oracle 10g ASM database. Although DGs can be created when creating an ASM instance in DBCA, DGs can also be created from SQL*Plus, which is the method shown in this application brief. This method is also efficient when creating DGs with multiple volumes. To create a DG, connect to an ASM instance and use the create disk group command. Designate the symbolic links referencing the appropriate HDLM devices.

It is recommended that you create all ASM disk groups with EXTERNAL Redundancy. The Universal Storage Platform has sufficient redundancy at the hardware level using any RAID level, so there is no need for ASM disk mirroring.

Configuring ASM Disk Groups and CRS Volumes

You will need a minimum of three ASM Disk Groups to work with TrueCopy or ShadowImage replication with Oracle Database 10g ASM in Hot Backup mode. The first two (DATADG and FLASHDG) will be replicated in separate consistency groups while the third (REDODG) is not replicated. In the examples outlined below, we created a DATADG disk group for data files, and a FLASHDG disk group for archive logs, control files, and flashback area. Both of these disk groups are replicated. We also created a nonreplicated REDODG disk group for redo logs. This configuration is shown in Table 1 and used in Hot Backup and Recovery procedures as demonstrated in Scenarios 1 and 2. If RAC is used, then you would also need at least one other LUN for CRS to include the OCR disk and voting disk. Note: the Cluster Ready Services (CRS) files (OCR/Voting disk) cannot be stored in ASM. The CRS volumes need to be identical on each instance.

You may decide to create multiple ASM disk groups to separate the Flash Recovery Area and Archive Logs, and have multiple DATA disk groups (such as separating read versus write activity tablespaces). A 5DG + CRS configuration is shown in Table 2. The procedures for Backup and Recovery would require corresponding ShadowImage group changes, using the correct svrctl commands to handle the RAC instances, and adding procedures to back up the CRS files, which include the OCR and Voting Disks.

For the OCR disk, a backup is done automatically once every four hours by default to the \$ORA_CRS/HOME/cdata/crs directory, so a manual backup is not normally required.

A Voting disk backup should be performed manually as it is not scheduled automatically. Verify that the Voting disk device name is correct. You may choose to use a different symbolic link name.

```
$ dd if= /dev/asm/lu999 of=bkup.voting
```

Tablespaces

Tablespaces are located in Data Disk Groups. You may choose to create more than one DATA disk group for specific operational or performance reasons. Most often a single DATA disk group will be preferred. This simplifies the backup procedures and spreads the data evenly over the available resources.

Initialization Parameter Files

Use server parameter files for initialization parameter files. Server parameter files are located in the Data Disk Group area. Server parameter files are located here because they can be backed up using ShadowImage, TCA, or HUR software storage-based replication.

Control Files

Control files are also located in a Data Disk Group area and another copy in the Flash Disk Group to follow Maximum Available Architecture (MAA) guidelines.

REDO Log Files

To provide the greatest possible performance, REDO log files are located in a dedicated REDO Log Disk Group. Because REDO log files do not need to be backed up, you can avoid unnecessary backup by separating these files from other DGs.

REDO log files are made up of two members/three groups for each node. Although you can use a one-member configuration, we recommend a two-member configuration.

Archive Log Files

It is essential to back up archive log files during hot backups to convert large-size REDO log files that are created when there is a hot backup into archive log files. However, since backups must be made according to a different schedule from data files, a dedicated DG is required.

Archive log files are written sequentially, so the use of storage systems (such as Hitachi Lightning 9900™ V Series or Hitachi Adaptable Modular Storage systems) that are lower in capacity unit cost than RAID-5 or the USP models may be considered, especially if they are virtualized behind a USP model.

Flashback Recovery Area

There may be a loss of the most recent archive log files if you use ShadowImage, TCA, or HUR software to restore the Archive log file disk group. To avoid this, we have shown the alternate configuration with additional disk groups in Scenarios 3 and 4. You can back up the archive log files in the Flashback Recovery Area DG (FLASHDG) before restoring the Archive Log DG (ARCHDG) with ShadowImage, TCA, or HUR software.

Layout examples for DGs are illustrated in Tables 1 and 2.

Table 1. Configuration Layout Using Three-Disk Groups in a Single Instance Environment

ASM Disk Group	Application
DATADG	SYSTEM table region SYSAUX table region UNDO table region USERS table region TEMP table region SPFILE Other User tablespaces
REDODG	Thread 1 Member 1 REDO log file 1 Thread 1 Member 1 REDO log file 2 Thread 1 Member 1 REDO log file 3 Thread 2 Member 1 REDO log file 1 Thread 2 Member 1 REDO log file 2 Thread 2 Member 1 REDO log file 3 Control file 1
FLASHDG	Archive logs Flashback Recovery Area Control file 2

Table 2. Configuration Layout Using Five-Disk Groups in a Two-node RAC Environment

ASM Disk Group	Application
DATADG1	SYSTEM table region SYSAUX table region UNDO table region USERS table region TEMP table region SPFILE Other USER tablespaces
DATADG2	Other USER tablespaces
REDODG	Thread 1 Member 1 REDO log file 1 Thread 1 Member 1 REDO log file 2 Thread 1 Member 1 REDO log file 3 Thread 1 Member 2 REDO log file 1 Thread 1 Member 2 REDO log file 2 Thread 1 Member 2 REDO log file 3 Thread 2 Member 1 REDO log file 1 Thread 2 Member 1 REDO log file 2 Thread 2 Member 1 REDO log file 3 Thread 2 Member 2 REDO log file 1 Thread 2 Member 2 REDO log file 2 Thread 2 Member 2 REDO log file 3 Control file 1
ARCHDG	Archive logs
FLASHDG	Flashback Recovery Area Control file 2

Using ASM Instances

Starting and Stopping ASM Instances

ASM instances are booted and stopped from CRS using srvctl commands. You can execute from any node, but can only control that one node. You cannot control multiple nodes at one time.

Starting multiple ASM instances:

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

Stopping multiple ASM instances:

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

Moreover, in the case of a single instance, starting and stopping is performed by setting +ASM in the SID environment variable, the same as in normal database instances.

Starting a 10g ASM instance:

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

Stopping a 10g ASM instance:

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

Listing the ASM Instance Status

You can find out if an ASM instance has started by using crs_stat command. (See Table 3.)

```
$ crs_stat -t Name Type Target State Host
```

Table 3. ASM Instance Status

ora....DB1.srv	application	ONLINE	ONLINE	nodename1
ora....DB2.srv	application	ONLINE	ONLINE	nodename2
ora....TPCC.cs	application	ONLINE	ONLINE	nodename1
ora....B1.inst	application	ONLINE	ONLINE	nodename1
ora....B2.inst	application	ONLINE	ONLINE	nodename2
ora.USPDB.db	application	ONLINE	ONLINE	nodename1
ora....SM1.asm	application	ONLINE	ONLINE	nodename1
ora....7C.lsnr	application	ONLINE	ONLINE	nodename1
ora....47c.gsd	application	ONLINE	ONLINE	nodename1
ora....47c.ons	application	ONLINE	ONLINE	nodename1
ora....47c.vip	application	ONLINE	ONLINE	nodename1
ora....SM2.asm	application	ONLINE	ONLINE	nodename2
ora....7D.lsnr	application	ONLINE	ONLINE	nodename2
ora....47d.gsd	application	ONLINE	ONLINE	nodename2
ora....47d.ons	application	ONLINE	ONLINE	nodename2
ora....47d.vip	application	ONLINE	ONLINE	nodename2

Using a method identical to normal databases, you can find out the status of an ASM instance.

```
SQL> select instance_name, status from v$instance;
```

However, an ASM instance does not become OPEN. The normal operating status is MOUNT.

Using ASM Disk Groups

As a handy summary, this section summarizes the ASM Disk Group commands most often used in setup.

Adding and Deleting Disks from DGs

With ASM, you can add and delete disks from disk groups online, even when a database instance is running.

To add disk /dev/sddlmae to disk group DATADG use:

```
SQL> alter disk group DATADG add disk '/dev/sddlmae';
```

For disk group deletions, first find out the disk name from v\$asm_disk view:

```
SQL> select path, name, state from v$asm_disk;
```

For instance, when the disk name is DATADG_0099 corresponding to /dev/sddlmba, you can delete it as follows:

```
SQL> alter disk group DATADG drop disk DATADG_0099;
```

DG Information Acquisition Method

One of the extremely important tasks in database administration is to check the disk group free area and mount status.

You can obtain information about the ASM disk group status from the V\$ASM_DISK GROUP dictionary view.

How to check the disk group mount status:

```
$ export ORACLE_SID=+ASM
sqlplus "/ as sysdba"
SQL> select group_number, name, state from v$asm_disk;
```

In the unlikely event that a DG is not mounted, you can mount it:

```
SQL> alter disk group DATADG mount;
```

To check the degree of redundancy of a DG, use:

```
SQL> select name, type from v$asm_disk;
```

To check DG capacity and amount of free space, use:

```
SQL> select name, total_mb, free_mb from v$asm_disk;
```

ASM Files

Under most operating conditions, there is no need to operate ASM files directly. However, you may have to operate on an ASM file directly for particular reasons. Moreover, when you cannot directly specify an Oracle-Managed File (OMF) file like a CREATE CONTROLFILE, you may have to operate an ASM file directly when you create an alias, as well.

Deleting ASM Files

Because ASM files cannot be moved or copied, they must be deleted and then recreated. To delete an ASM file:

```
$ export ORACLE_SID=+ASM  
sqlplus "/ as sysdba"  
SQL> alter disk group DATADG drop file  
'+DATADG/db10asm/tablefile/controlfile.256.1' ;
```

Creating and Deleting Aliases and Changing Names

All ASM files are provided with a group_number, file_number and file name. These attributes cannot be changed because they are controlled by OMF. To handle this, you can attach only one alias for files in ASM. You can attach the names you wish to files managed by OMF with this feature, thereby enhancing manageability.

For example, you can provide the alias +ARCHDG/controlfile_1 to a file created by OMF and called +ARCHDG/db10asm/controlfile/current_1.256.1:

```
$ export ORACLE_SID=+ASM  
sqlplus "/ as sysdba"  
SQL> alter disk group DATADG add alias '+DATADG/controlfile_1' for  
'+DATADG/db10asm/controlfile/current_1.256.1';
```

If there is already an alias that has been created, it will result in an error, so change the existing alias name:

```
$ export ORACLE_SID=+ASM  
sqlplus "/ as sysdba"  
SQL> alter disk group DATADG rename alias '+DATADG/old_name' to  
'+DATADG/new_name' ;
```

Aliases are deleted as follows:

```
$ export ORACLE_SID=+ASM  
sqlplus "/ as sysdba"  
SQL> alter disk group DATADG drop alias '+DATADG/controlfile_1';
```

ASM File Information Acquisition Method

Use dictionary view or Enterprise Manager to acquire ASM file information, rather than using a shell.

When you want to check the file name and date created or date modified:

```
SQL> select name, to_char(creation_date,'yyyy-mm-dd hh24:mi:ss') creation_date,
  to_char(modification_date,'yyyy-mm-dd hh24:mi:ss') modification_date, from
  v$asm_alias a, v$asm_file b where a.file_number=b.file_number and
  a.group_number=b.group_number
  order by name ;
```

When you want to check the file capacity, free space, type (data file, control file, etc.):

```
SQL> select name, bytes, space, type from v$asm_alias a, v$asm_file b where
  a.file_number=b.file_number and a.group_number=b.group_number order by name ;
```

When you want to check file redundancy and stripe granularity:

```
SQL> select name, redundancy, striped from v$asm_alias a, v$asm_file b where
  a.file_number=b.file_number and a.group_number=b.group_number order by name ;
```

ShadowImage Software Recommended Settings

Pair Configuration

Backups can be performed in ShadowImage groups. These ShadowImage groups should contain all the same volumes as referenced in the corresponding ASM disk groups. For example, with the minimum three ASM disk groups method discussed earlier, the first two (DATADG and FLASHDG) will be replicated in separate consistency groups while the third (REDODG) is not replicated. These ShadowImage groups need to be updated when new LUNs are added or removed from the corresponding ASM disk group.

Pair Control

When pairs are created, “Point-in-Time (PiT) split” options must be provided for each consistency group that corresponds to a given ASM disk group.

The command line method of controlling ShadowImage or TrueCopy pairs is to use Hitachi Command Control Interface (CCI). This software is installed on the hosts that manage the pairs. A special LUN called a command device is used to communicate with the storage system to perform the actual storage replication and to receive back status messages. We have installed CCI on the primary host in our scenarios described later in this paper. HORCM#.conf configuration files shown in Appendix C to define the source and target LUNs for the ShadowImage pairs. CCI Instance 0 is defined by HORCM0.conf. CCI Instance 1 is defined by HORCM1.conf and describes the LUNs that will be mounted on the backup host.

Refer to the Hitachi Command and Control Interface (CCI) manual for more detailed information on installing, configuring, and using CCI commands.

These CCI instances are started before issuing ShadowImage CCI commands by issuing:

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

Once the HORCM#.conf configuration files are defined and the instances are started CCI commands can be issued to establish the ShadowImage pairs. The syntax is below. The -IM0 flag indicates that we are using instance 0 as the local instance and treating the command as (M) ShadowImage. TrueCopy would be indicated with an H instead of the M and have additional flags. The -vl flag establishes the direction of the copy (vector local), which means (local instance to remote instance). The best practice is to always communicate with the CCI instance that contains the primary LUNs and establish the pairs in the normal direction using the -vl flag. The -m grp option indicates to use ShadowImage group consistency.

Here is an example of ShadowImage consistency group creation:

```
# paircreate -IM0 -g DATADG -vl -m grp  
# paircreate -IM0 -g FLASHDG -vl -m grp
```

Typically, several volumes are contained in a disk group. By attaching a PiT option, consistency can be maintained between volumes so that it is as if you are backing up a single volume. If you create a pair without providing PiT options, you cannot be assured of maintaining consistency between volumes because the splits are not atomic across the entire group. In the worst-case scenario, disk groups are backed up in an improper status. The group is placed in suspend status after the initial copy or resynchronization is complete.

Here is an example of ShadowImage consistent split.

```
# pairsplit -IM0 -g DATADG  
# pairsplit -IM0 -g FLASHDG
```

In the following example of ShadowImage wait check for status of psus, the -s flag can check for a variety of pair states (smpl, pair, copy, psus, etc.) and the -t flag waits up to the indicated time in seconds. If the status waited for completes before this time, then the command returns earlier than the full wait time. A production-level script should incorporate full error checking and verify all command return codes.

```
# pairevtwait -IM0 -g DATADG -s psus -t 9999  
# pairevtwait -IM0 -g FLASHDG -s psus -t 9999
```

Command Descriptions

Table 4 describes basic CCI pair commands and their functions.

Table 4. Pair Commands

Command Name	Description	Notes
paircreate	Establishes the pairs	
pairsplit	Temporarily splits pairs and suspends updating to the secondary volume; the pair status is maintained in order to continue holding the differential information	Suspends pairs
pairresync	Resynchronizes pairs that are in suspend status; can be returned to pair status at high speed, depending on the differential information	Command cannot be used when there is no diff information
pairresync restore	Resynchronizes pairs in the opposite direction when they are in suspending status and generates backup data lists; can be returned to pair status at high speed, depending on the differential information	Typically, a pair suspend is done after a restore pair list
pairsplit -S	Discards the pair differential information and deletes the pair; to put it in pair status again, you must recreate the pair	Deletes pair
pairevtwait	Command to wait until it becomes the specified pair; needed to make a positive transition to pair status	Pair status confirmation command
pairdisplay	Enables viewing of pair status and pair progress rate	Displays pair status
pairvolchk	You can check on whether or not there is a PIT option (Consistency Group setting), etc.	Pair attribute confirmation command

ASM Backup Using Hitachi ShadowImage, TrueCopy, or Universal Replicator Software

Backup Overview

Several basic points must be kept in mind in terms of backing up Oracle Database 10g. The following explanation is made in comparison with Oracle9i Database.

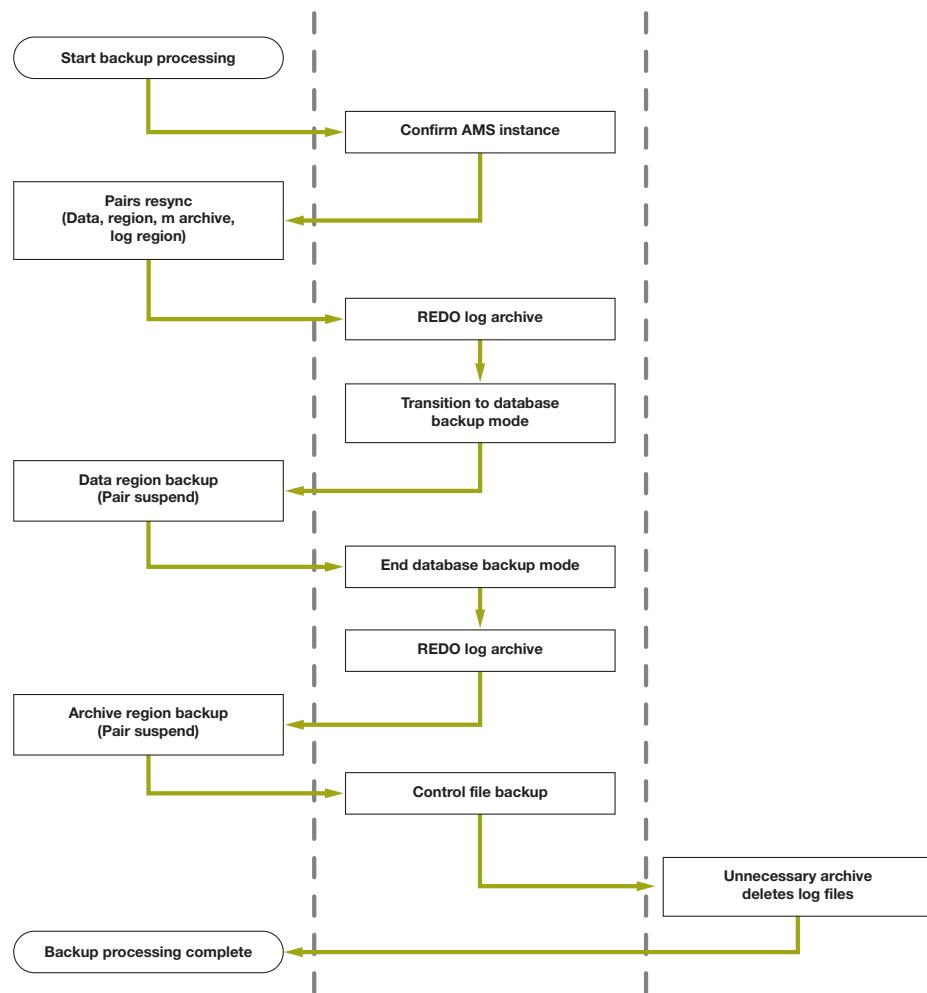
Although it is important to perform backups periodically, backups should be performed when there is a change in the database configuration or immediately after a recovery.

Comparing Oracle Database 10g and Oracle9i Database

The backup procedures for Oracle Database 10g are not much different from those for Oracle9i Database; the main difference is that consistency groups are required for ShadowImage and TrueCopy software. In Oracle Database 10g the ASM instances and the database have to be stopped when there is a cold backup of the database.

As shown in Figure 4, Oracle Database 10g ASM hot backups on Hitachi storage require the use of ShadowImage, TrueCopy, or HUR consistency groups to guarantee that atomic splits of all the volumes in a disk group occur at the same time.

Figure 4. Hot Backup Processing Flowchart



Cold Backup Precautions

When performing a cold backup, stop the database and the ASM instance.

Hot Backup Precautions

When performing a hot backup, provide PiT options when creating pairs. Hot backups in ASM must have PiT options attached to the pairs. If a pair suspend and backup is performed without PiT options provided, the volumes within a disk group will not be consistent with one another; there is a risk that invalid data could be incorporated. As a result, the hot backup could not be an accurate copy of the production database.

Backup Requirements for Specific Data Types

Tablespaces

As noted before in Oracle9i, tablespaces can be backed up by placing each tablespace in backup mode. This becomes a maintenance issue in that scripts need to be updated to check for new tablespaces as part of the Hot Backup procedures. Starting with Oracle 10g, the entire database can now be placed in backup mode with a single command, which simplifies backup operations. This is now the recommended best practice method.

Begin backup mode for database:

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

End backup mode for database:

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

Control Files

Control files are backed up using Oracle Recovery Manager ("RMAN"). For a handy summary of the most commonly used RMAN commands, please refer to Appendix A. The commands are:

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

Control files must be backed up with Oracle commands for recovery as the LUNs have already been split when these must be backed up.

Server Parameter Files

In order that they may be backed up by ShadowImage software, server parameter files can be backed up according to the same procedures as table regions. Furthermore, since server parameter files are located in DATADG, they are backed up along with DATADG when the DATADG is backed up.

REDO Log File

REDO log files cannot be backed up according to previous procedures. However, if you wish to back up the REDO log file contents, use alter system archive log current and output to archive log files. This method is possible because archive log files can be backed up as ASM files; you can create a REDO log file at time of recovery if the backup is performed with ShadowImage software. The contents cannot be used.

Archive Log Files

Although archive log files may have been output in the conventional file system, by backing them up by DG, backup is enabled because they were only output on the DG. When you do a hot backup, logs that are needed for media recovery are stored as REDO log files, so you must not forget to execute the appropriate Oracle ALTER SYSTEM commands (for brevity's sake, not summarized here) to archive log current and output to the archive log file.

Also, since the archive log file does not have a backup mode, performing the backup during archive log output can cause the destruction of files. Therefore, with ShadowImage software, you must execute the appropriate ALTER SYSTEM commands to archive log current when you perform a backup, confirm that the archive log output of all nodes has been positively completed, and then perform a pair split (backup). Confirm that archive log output has been completed with the alert log file of each node's DB instance. After the backup has been completed, you may delete the old archive log files because they are no longer required.

Deleting Old Archive Logs

Connect to Recovery Manager (RMAN). For a handy summary of the most commonly used RMAN commands, please refer to Appendix A:

```
$ export ORACLE_SID=db10asm  
$ rman target / nocatalog
```

Delete all archive log files that have been created more than two weeks previously:

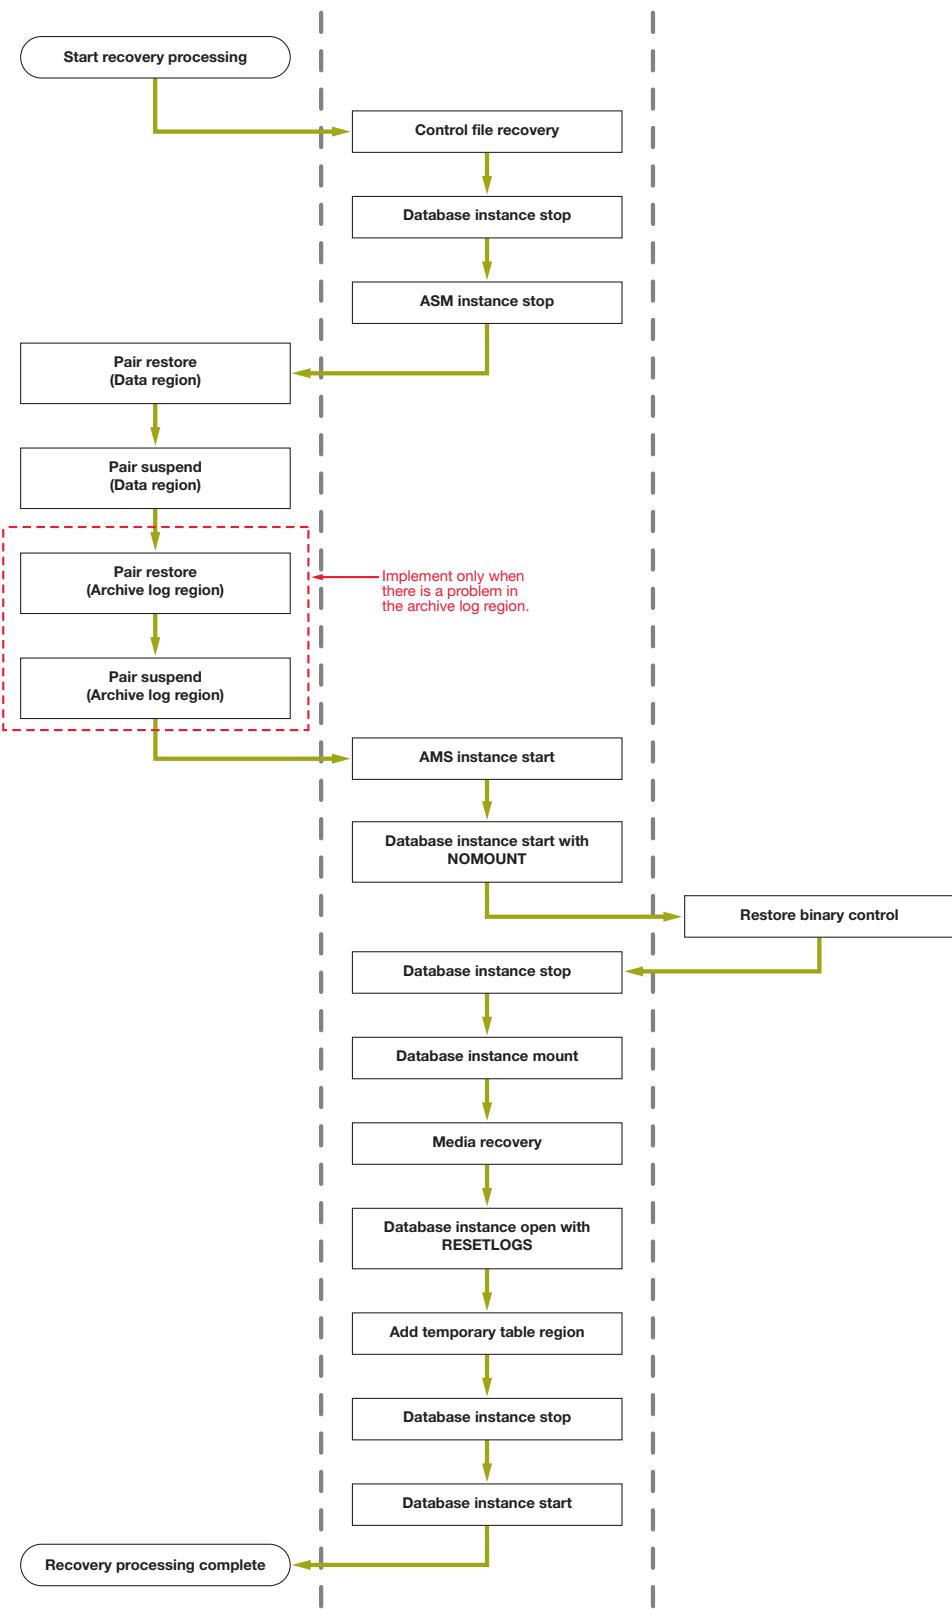
```
RMAN> DELETE ARCHIVELOG ALL COMPLETED BEFORE 'SYSDATE-14';
```

Recovery Overview

Recovery Procedure Flowchart

Figure 5 illustrates recovery procedures.

Figure 5. Recovery Procedure Flowchart



Precautions: Since the data group will be destroyed if you perform a restore while an ASM instance is running, be absolutely sure to perform restores after stopping ASM instances.

Scenario 1: Backup of Oracle Database 10g ASM Using ShadowImage Software

On Production Host

This procedure places the Oracle Database 10g ASM in Hot Backup mode on the production host for a short time while a consistent split copy is made with ShadowImage software. The copy can then be used on the backup server for testing, backup processing, or other tasks without further impact to the production server or database. Start with pairs in PAIR state. Create or resynchronize the pairs if needed. This example is using the minimum 3DG environment with DATADG, FLASHDG, and REDODG.

Establish the ShadowImage pairs or resynchronize if already created but in suspend state as needed:

```
# paircreate -IM0 -g DATADG -vl -m grp  
# paircreate -IM0 -g FLASHDG -vl -m grp  
or  
  
# pairresync -IM0 -g DATADG  
# pairresync -IM0 -g FLASHDG
```

Verify pairs are in paired (PAIR) state:

```
# pairevtwait -IM0 -g DATADG -s pair-t 9999  
# pairevtwait -IM0 -g FLASHDG -s pair -t 9999
```

Place Oracle Database 10g in Hot Backup mode on the production host:

```
$ export ORACLE_SID=db10asm  
$ sqlplus "/as sysdba"  
SQL> alter database begin backup;
```

Split the pairs; pairs will be split consistently because they were created using consistency groups:

```
# pairsplit -IM0 -g DATADG
```

Verify pairs are in a suspended (PSUS) state:

```
# pairevtwait -IM0 -g DATADG -s psus -t 9999
```

End backup mode for the Oracle Database 10g:

```
$ export ORACLE_SID=db10asm  
$ sqlplus "/as sysdba"  
SQL> alter database end backup;
```

Switch logs:

```
$ export ORACLE_SID=db10asm
$ sqlplus "/as sysdba"
SQL> alter system archive log current;
```

Create two Controlfile backups. One will be used to start the database (control_start) on the backup server in mount mode which as a side effect modifies the controlfile so another unmodified copy is needed. The other copy (control_backup) will be used by RMAN as the valid unmodified copy. For a handy summary of the most commonly used RMAN commands, please refer to Appendix A.

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

Resynchronize the RMAN catalog so that the latest archive log is known to RMAN:

```
RMAN> resync catalog;
```

Split the FLASHDG group which now has the latest archive log:

```
# pairsplit -IH0 -g FLASHDG
```

Verify pairs are in a suspended (PSUS) state:

```
# pairevtwait -IH0 -g FLASHDG -s psus -t 9999
```

Delete Unneeded Archive Log Files if Desired

This action deletes unneeded archive log files. Here, all archive log files that have been created more than two weeks previously are deleted using the first syntax. An alternate method is to delete archive logs backed up twice as shown in the second syntax. You may choose to keep archive logs longer if needed for flashback.

```
$ export ORACLE_SID=db10asm
$ rman target / nocatalog
RMAN> delete archivelog all completed before 'SYSDATE-14';
or

RMAN> delete archivelog backed up 2 times to device type 'sbt';
```

Processing is now complete on the production host.

On Backup Host

The backup host can now use the replicated volumes.

Start the ASM instance: The ASM instance will find the ShadowImage or TrueCopy S-vols even if the raw device or HDLM names are renumbered and different from the production host. The ASM_DISKSTRING will identify the volumes to search for as ASM volumes. Oracle Database 10g can identify information from these candidate volumes the proper disk group to which they belong. The ASM_DISKGROUPS parameter identifies those disk groups that will be mounted when ASM starts.

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

Start up the copied database on the backup server: This snapshot (from the S-VOLs) can be mounted in a variety of modes depending on the purpose intended. You can mount it to allow RMAN to back up the database or keep it as a quick recovery option to restore back to the primary host using recovery options (such as resetlogs, roll-forward recovery with additional logs, etc.). An RMAN backup is described below. A recovery back to the primary host is described in Scenario 2. Scenario 5 shows a cloning technique that resets the logs on a secondary host.

The following example uses RMAN to back up the database on the backup server. The startup controlfile is used to mount the database, then the backup controlfile is used as part of the RMAN backup.

Change the backup server init.ora file to point to the startup controlfile:

```
init.ora file  
  
control_files = +FLASHDG/control_file/control_start
```

Mount the database instance on the backup server:

```
$ export ORACLE_SID=db10asm  
$ sqlplus "/as sysdba"  
SQL> startup mount
```

Perform backup using RMAN:

```
RMAN> run {  
    allocate channel t1 type sbt_tape  
    backup format 'ctl %d/%s/%p/%t'  
    controlfilecopy '+FLASHDG/control_file/control_backup';  
    backup  
        full  
        format 'al %d/%s/%p/%t'  
        (archive all);  
    release channel t1  
}
```

Scenario 2: Recovery of Oracle Database 10g ASM Using ShadowImage Software

ShadowImage software can provide a quick recovery method of restoring data files taken during a previous Hot Backup. Downtime is minimized due to the fact that only changed tracks based on differential bitmaps are required to be restored back to the primary LUNs. Once data files are restored then archive logs since the time of the restored backup may be applied for database instance recovery. Online redo logs in the REDODG disk group and archive logs in the FLASHDG disk group are not normally restored unless they are also damaged or otherwise invalid. The REDODG contains the last committed transactions. If they are invalid, then recreate the disk group and open the database with resetlogs option. If the FLASHDG disk group is invalid, then use ShadowImage software to restore the disk group.

Shut down the database instance on the production server if it is not already down:

```
$ export ORACLE_SID=db10asm  
$ sqlplus "/as sysdba"  
SQL> shutdown immediate
```

Shut down the ASM instance or dismount the ASM datafiles that need to be restored if the ASM instance is needed for other databases:

```
$ export ORACLE_SID=+ASM  
$ sqlplus "/as sysdba"  
SQL> shutdown immediate  
or  
SQL> alter diskgroup DATADG dismount;
```

ShadowImage DATADG disk group should be in suspended (PSUS) state already:

```
# pairresync -IM0 -g DATADG -restore
```

Verify pairs are in a PAIR state:

```
# pairevtwait -IM0 -g DATADG -s pair -t 9999
```

Suspend the group again so that the backup copy is unchanged and available for recovery again if needed:

```
# pairsplit -IM0 -g DATADG
```

Verify pairs are in a suspended (PSUS) state:

```
# pairevtwait -IM0 -g DATADG -s psus -t 9999
```

Start or mount the ASM disk group as required:

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

or

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

Start up the restored production database mount only.

Choose whether you want to perform complete or PiT recovery using RMAN.

PiT recovery can be specified with the “until SCN” or “until time” options:

```
$ RMAN> run {
          recover database;
      }
```

or

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

Open the database:

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

Scenario 3: Oracle Database 10g ASM Cloning Using ShadowImage Software

ShadowImage software commands are shown but this can also be used with TrueCopy and HUR software or TrueCopy and HUR software with a remote ShadowImage cascaded copy.

This procedure creates a PiT copy of the Oracle Database 10g ASM without placing the database in Hot Backup mode. The copy can then be used on another server for testing purposes, but cannot be used for recovery of the production database. The ShadowImage group needs to include the data files, controlfiles, redo logs, and flashback logs. The archive logs are not required; however, since our controlfiles and flashback area are all in the FLASHDG group, we have included all LUNs in a single ShadowImage consistency group that will be part of a consistent At-Time split.

Start with pairs in PAIR state. Create or resynchronize the pairs if required:

```
# paircreate -IM0 -g SIALLDB -vl
or
```

```
# pairresync -IM0 -g SIALLDB
```

Split the pairs. Pairs will be split consistently because they were created using consistency groups:

```
# pairsplit -IM0 -g SIALLDB
```

Verify pairs are in a suspended (PSUS) state:

```
# pairevtwait -IM0 -g SIALLDB -s psus -t 9999
```

Start the ASM instance on the test server:

Verify that the ShadowImage Svos on the test server have the correct Oracle permissions and that the ASM_DISKSTRING and ASM_DISKGROUPS are specified correctly:

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

Mount the cloned database on the test server:

```
$ export ORACLE_SID=db10asm
$ sqlplus "/as sysdba"
SQL> startup mount
SQL> recover database;
SQL> exit
```

The database can optionally be renamed using the Oracle nid utility at this point:

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

Open the cloned database with resetlogs:

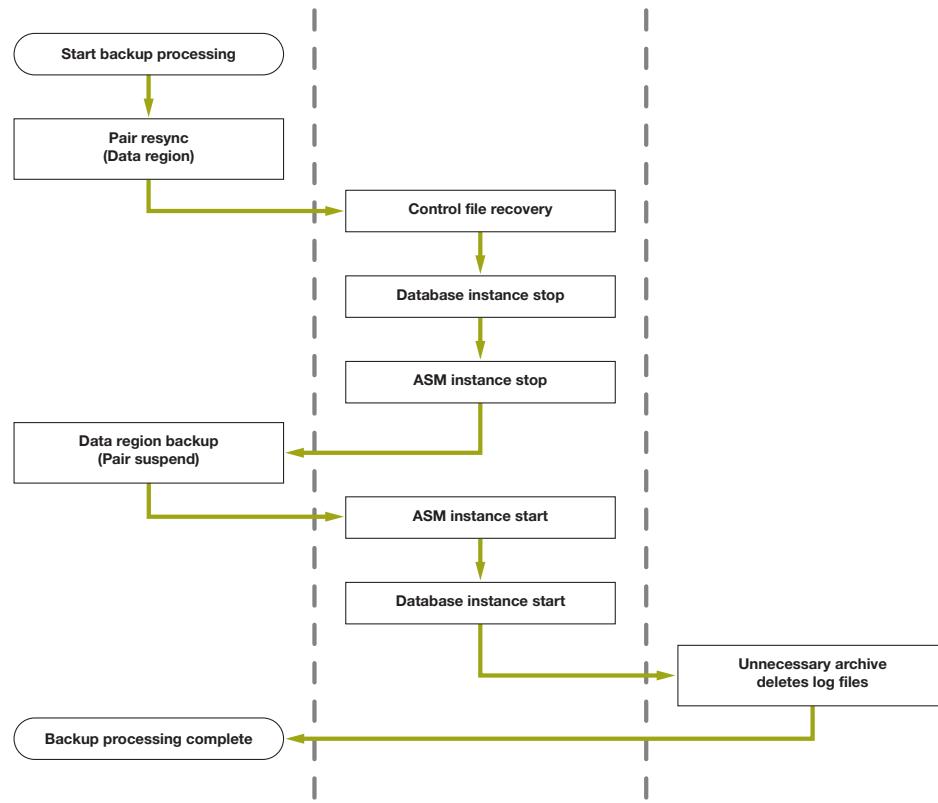
```
$ export ORACLE_SID=oratest
$ sqlplus "/as sysdba"
SQL> startup mount
SQL> alter database open resetlogs;
```

The alternate host can now use the cloned copy of the database. The database will be in a state as if the server had a power failure and the database will go through normal crash restart recovery. Once this is complete the database will be a usable cloned copy.

Scenario 4: Oracle Database 10g Cold Backup Using ShadowImage Software's Overall Flowchart of Backup Procedures

Cold backup procedures for Oracle Database 10g ASM are identical to procedures for Oracle Database 10g non-ASM and Oracle9*i*, with one key exception, which is to also stop all ASM instances first in addition to the database instance before taking a cold backup. (See Figure 6.)

Figure 6. Cold Backup Processing Flowchart



ShadowImage commands are shown, but this can also be used with TrueCopy and HUR software or TrueCopy and HUR with a remote ShadowImage cascaded copy.

This procedure creates a copy of the production database while the database is shut down. No I/O is occurring so this is a straightforward volume copy process. Consistency groups are not required. The copy can then be used on another server for testing purposes, backed up with RMAN on a remote server, or used for quick recovery by restoring back to the production server. The ShadowImage group needs to include the data files, controlfiles, redo logs, and flashback logs. The archive logs are not required; however, since our controlfiles and flashback area are all in the FLASHDG group, we have included all LUNs in a single ShadowImage consistency group.

Start with pairs in PAIR state. Create or resynchronize the pairs if required:

```
# paircreate -IM0 -g SIALLDB -vl
or
# pairresync -IM0 -g SIALLDB
```

Stop the ASM instance if it is running. The database should already be shut down:

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

Split the pairs. Pairs will be split consistently because they were created using consistency groups:

```
# pairsplit -IM0 -g SIALLDB
```

Verify pairs are in a suspended (PSUS) state:

```
# pairevtwait -IM0 -g SIALLDB -s psus -t 9999
```

The ShadowImage copy is now a valid copy of the database that can be immediately used for recovery by restoring back to the primary LUNs or mounted for testing purposes on a secondary server.

Appendix A: RMAN Recovery Catalog Guidelines

For brevity, a summary of most commonly used commands and some best practices for Oracle Recovery Manager (“RMAN”) are summarized in this Appendix.

Always connect to the same recovery catalog, whether on the production server or the backup server.

Connect to the recovery catalog with the production database as the target on the production host:

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

Connect to the recovery catalog with the backup database as the target on the backup host:

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

The following is an example of creating a recovery catalog:

```
$ rman catalog rman/rman@rcat
RMAN> create catalog;
RMAN> exit

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

Target Database Password

Connected to target database:

```
DB10ASM (DBID=3974439867)
```

Connected to recovery catalog database

```
RMAN> register database;
```

Database registered in recovery catalog

Starting full resync of recovery catalog

Full resync complete

```
RMAN> report schema;
```

Report of database schema

List of Permanent Datafiles

	File	Size(MB)	Tablespace	RB	segs	Datafile Name
1	490		SYSTEM	YES		+DATADG/db10asm/datafile/system.256.617205169
2	25		UNDOTBS1	YES		+DATADG/db10asm/datafile/undotbs1.258.617205171
3	310		SYSAUX	NO		+DATADG/db10asm/datafile/sysaux.257.617205169
4	5		USERS	NO		+DATADG/db10asm/datafile/users.259.617205171
5	100		EXAMPLE	NO		+DATADG/db10asm/datafile/example.261.617205255

List of Temporary Files

	File	Size(MB)	Tablespace	Maxsize(MB)	Tempfile	Name
1		20	TEMP	32767		+DATADG/db10asm/tempfile/temp.260.617205251

Appendix B: Sample TNS Configuration File

This is a sample Oracle TNSNAMES.ora configuration file.

```
# -----
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) ))
```

Appendix C: Sample ShadowImage HORCM Configuration Files

This are samples of USP HORCM configuration files. The recommended file directories and names are shown as well as the files' contents.

/etc/horcm0.conf

```
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
172.17.173.83   11000        1000          3000

HORCM_CMD
#dev_name        dev_name     dev_name
/dev/sdx

HORCM_LDEV
#dev_group       dev_name     Serial#      CU:LDEV(LDEV#) MU#
DATADG          0070_0150   11210        00:70      0
DATADG          0071_0151   11210        00:71      0
DATADG          0072_0152   11210        00:72      0

FLASHDG         0073_0153   11210        00:73      0
FLASHDG         0074_0154   11210        00:74      0

SIALLDG         0070_0230   11210        00:70      1
SIALLDG         0071_0231   11210        00:71      1
SIALLDG         0072_0232   11210        00:72      1
SIALLDG         0073_0233   11210        00:73      1
SIALLDG         0074_0234   11210        00:74      1
SIALLDG         0080_0240   11210        00:80      1
SIALLDG         0081_0241   11210        00:81      1

HORCM_INST
#dev_group      ip_address   service
DATADG          172.17.173.83 11001
FLASHDG         172.17.173.83 11001

SIALLDG         172.17.173.83 11001
```

/etc/horcm1.conf

```
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
172.17.173.83   11001        1000          3000

HORCM_CMD
#dev_name        dev_name     dev_name
/dev/sdx

HORCM_LDEV
#dev_group       dev_name     Serial#      CU:LDEV(LDEV#) MU#
DATADG          0070_0150   11210        01:50      0
DATADG          0071_0151   11210        01:51      0
DATADG          0072_0152   11210        01:52      0

FLASHDG         0073_0153   11210        01:53      0
FLASHDG         0074_0154   11210        01:54      0
```

SIALLDG	0070_0230	11210	02:30	0
SIALLDG	0071_0231	11210	02:31	0
SIALLDG	0072_0232	11210	02:32	0
SIALLDG	0073_0233	11210	02:33	0
SIALLDG	0074_0234	11210	02:34	0
SIALLDG	0080_0240	11210	02:40	0
SIALLDG	0081_0241	11210	02:41	0

HORCM_INST				
#dev_group	ip_address	service		
DATADG	172.17.173.83	11000		
FLASHDG	172.17.173.83	11000		
SIALLDG	172.17.173.83	11000		

Appendix D: Sample Contents of ASM Instance Parameter File

For reference, this appendix describes key parameters included in the ASM Database Instance parameter file.

```
INSTANCE_TYPE=ASM
ASM_DISKSTRING='/dev/sdd1m*'
ASM_DISKGROUPS='DATADG, FLASHDG, REDODG'
```

Sample contents of Database Instance parameter file:

```
db_name=db10asm
control_files='+DATADG/control_001
DB_RECOVERY_FILE_DEST='+FLASHDG
LOG_ARCHIVE_DEST_1='LOCATION=USE_DB_RECOVERY_FILE_DEST'
```

Hitachi Data Systems Corporation

Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA

Contact Information: 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA

Contact Information: 1 408 970 1000 info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom

Contact Information: + 44 (0) 1753 618000 info.uk@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.

TrueCopy is a registered trademark and Universal Storage Platform, ShadowImage, Universal Star Network, Network Storage Controller, and Lightning 9900 are trademarks of Hitachi Data Systems Corporation.

IBM and FICON are registered trademarks of International Business Machines Corporation.

Microsoft is a registered trademark of Microsoft 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/products_services/support/warranty.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 2007. All Rights Reserved.

WHP-221-01 LWD June 2007