



Best Practices for Oracle Automatic Storage Management (ASM) on Hitachi Dynamic Provisioning(HDP)

Rev. 0

March 2008




Introduction

Until now, when storage was introduced, generally future capacity requirements were studied and at first the minimum required capacity was purchased and expansion disks added as necessary. However, in current IT systems, with the spread of Internet business, and video and movie content, the amount of data handled daily is rapidly increasing and since an increase in future storage requirements is forecast, initial introduction costs are rising and estimating and managing capacity is becoming extraordinarily complex.

When the volumes storing such increasing data (example: volumes for Oracle database) are secured as virtual volumes with Hitachi Dynamic Provisioning software (HDP) for the Hitachi Storage Universal platform V (USP V), costs can be reduced and the management burden reduced.

This document explains some best practices for using Hitachi Dynamic Provisioning software and Oracle Automatic Storage Management(ASM), which have been derived as a result of tests and analysis, performed in the Hitachi-Oracle evaluation centers.



Contents

Introduction	2
Hitachi Universal Storage Platform V (USP V).....	4
Hitachi Dynamic Provisioning software (HDP).....	5
■ Overview of Dynamic Provisioning	5
■ Primary Component.....	6
■ I/O processing.....	9
Oracle Automatic Storage Management (ASM).....	10
Oracle ASM on HDP	14
■ Improved initial costs of storage ownership and capacity design.....	14
■ Improvement of disk management for online disk addition.....	14
■ Improved capacity management cost.....	16
■ Improved performance design cost.....	16
Best practice of Oracle ASM with HDP	17
■ Approach to a system configuration design.....	17
■ Performance.....	21
■ Backup & restore.....	23
Important observation.....	31
Appendix A: USP-V Specifications.....	32
Appendix B: HDP Support Configuration Summary	33
Appendix C: RAID Manager Configuration Definition File Example.....	33
Appendix D: ShadowImage Pair Status Monitoring	34

Hitachi Universal Storage Platform V (USP V)

USP V not only improves the performance and scalability of the Hitachi Universal Storage Platform (USP), but is also the world's first disk sub-system that further evolves enterprise class virtualization functions and virtualizes volume capacity. It has become possible to actually sense the advantages of storage solutions equipped with the industry's highest level performance and expandability for a wide variety of customers from medium-size growth companies to large global corporations. Up to 128 disk drives can be mounted in one basic chassis with disk control section (DKC: Disk Controller) and up to 1152 disk drives can be mounted with adding four disk frame chassis (DKU: Disk Units).



Figure 1. Hitachi Storage Universal Platform V (USP V)

Hitachi Dynamic Provisioning software (HDP)

■ Overview of Dynamic Provisioning

Hitachi Dynamic Provisioning software (HDP) is a new program product that improves the overall efficiency of customer storage usage and reduces initial storage costs and management costs. Write access to virtual volumes is realized through the necessary minimum assignment of the pool volume made up of the actually mounted physical volumes.




Figure 2. Overview of HDP

■ Primary Component

Hitachi Dynamic Provisioning (HDP) uses four primary components:

- HDP Pool
- HDP Pool volumes
- Virtual volumes
- Virtual volume management area

HDP uses data on pool volumes via virtual volumes. Virtual volume is managed by Virtual volume management area in shared memory and is associated with a HDP pool. In order to use HDP, at least one virtual volume and one HDP pool are required.

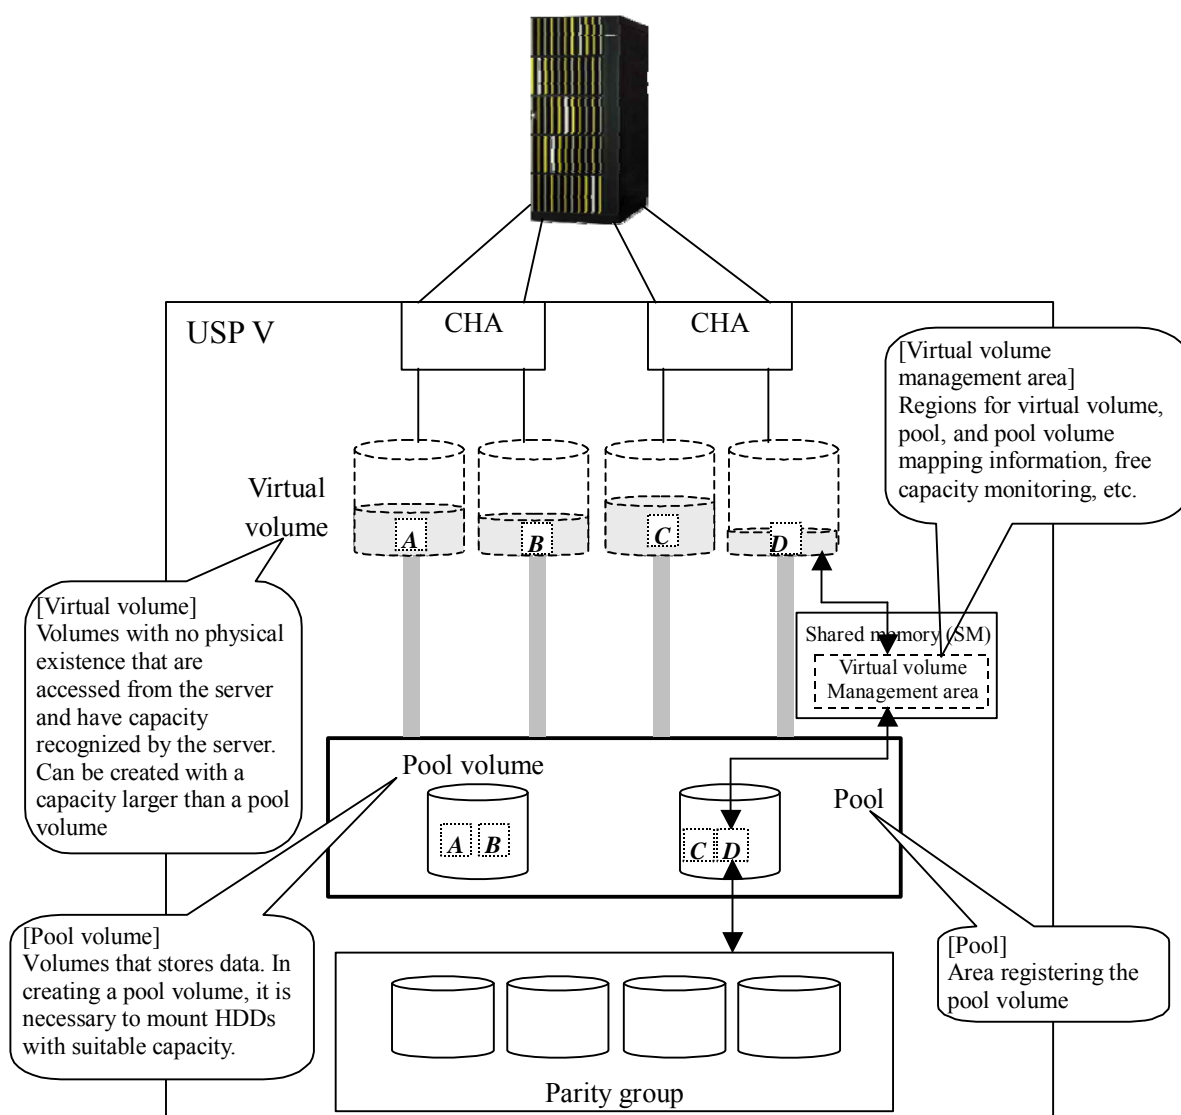



Figure 3. Primary component of HDP



■ Benefits

As users have adopted and deployed HDP for the volumes that are estimated to grow, they have seen the following benefits:

(1) Improved initial costs of storage ownership.

Dynamic Provisioning enables the movement to a just-in-time model of storage provisioning. i.e rather than mounting all storage capacity that will be required in the future at an early stage, they can be mounted as needed. This eliminates the waste from the storage infrastructure and reduces the initial and total cost of storage ownership.

(2) Improvement capacity design and management costs

To use HDP, initial capacity planning of virtual volumes includes future growth with initial HDP pool capacity reflecting immediate storage requirements. Initial dual planning of current and future capacity reduces the overall design and management cost normally associated with traditional methods.

Virtual volumes and available physical disk space are constantly monitored to notify administrators of relevant events and warnings.

(3) Improved performance design cost

Using wide striping techniques, Dynamic Provisioning automatically spreads the I/O load of all applications accessing the common pool across the available spindles. This process should reduce the chance of hot spots and optimizes I/O response times and can reduce the cost of storage performance design that usually requires the physical storage configuration design for performance.




Figure 4. Comparison with HDP and without HDP

■ I/O processing

Write I/O to a virtual volume of HDP is stored in the HDP pool volume, which is a volume with the actual pre-set capacity. Below are details of the I/O process when free space for HDP pool volume is not available for write access to a virtual volume from servers.

(1) A write command to a virtual volume is received from the server.

(2) Free space is assigned to a pool volume.

- The assigned free space is distributed to RAID groups so that the load is not concentrated on a specific parity group.

- The correspondence between the assigned free space and the virtual volume is managed in the virtual volume management area in shared memory of USP.

(3) The Write data is stored into cache memory and the I/O process is returned to the server.

(4) The Write data in the cache memory is written to the pool volume asynchronously.




Figure 5 HDP I/O Processing

Oracle Automatic Storage Management (ASM)

Automatic Storage Management (ASM) is a feature that was first introduced with Oracle Database 10g. ASM makes it possible to stripe and mirror data automatically and to add and delete disks online, and, since data is reallocated automatically, ASM makes physical file management easy. Introducing ASM makes it possible to greatly reduce the amount of work required for database administrators.

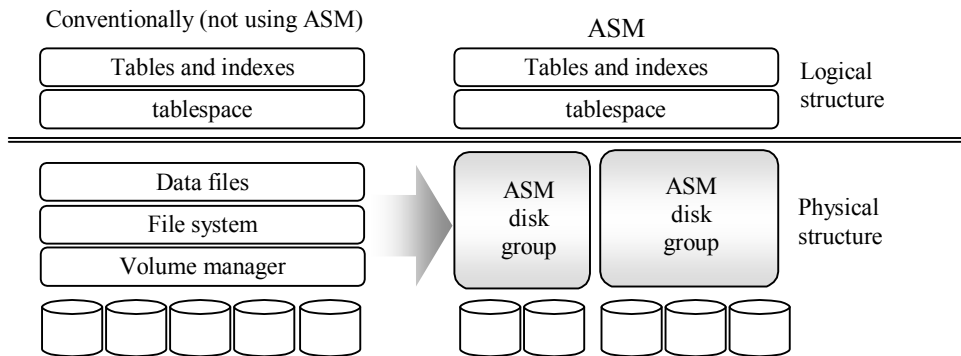


Figure 6. ASM Area Management Performing Storage Virtualization

(1) Disk group

ASM manages multiple disks as one logical unit, a disk group. Various database files - control files, data files, temporary files, redo files, and archive logs – are created in the disk group and distributed (striped) through it. The advantages of using a disk group are as follows:

- Sharing files among servers
- Adding/deleting disks online
- Storage can be configured without consideration for disk configuration or file layout.
- The storage method is the same, even for different platforms.

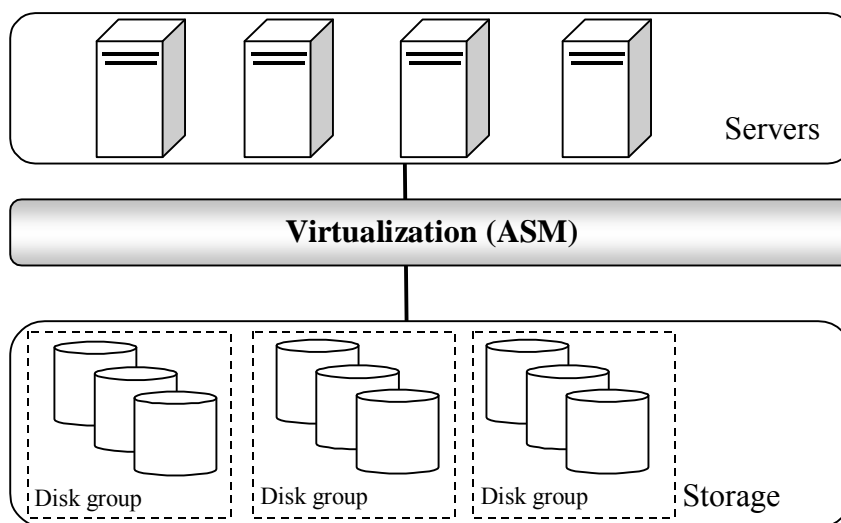


Figure 7. Disk Groups and ASM Virtualizing Disk Configuration

(2) ASM instance

Using ASM requires instance for storage management. This ASM instance must be started before startup of the database instance that will use the ASM. An ASM instance manages meta data for physical disks and provide layout information (maps) and mirror information for database instances.




Figure 8. Relationship of ASM Instances and Database Instances

(3) ASM files

Files written to disk groups are called ASM files. When an ASM file is created, the file name includes the "+<ASM disk group name>" path. These can not be accessed from normal OS interfaces.

(4) Striping

ASM distributes data in such a way as to even out the usage rates for all the disks in the disk group. This evens out access to the disk as a whole and prevents “hot spots,” meaning specific locations where disk access concentrates.




Figure 9. ASM Striping Function

With Oracle 10g, striping areas are allocated in allocation units (AU) of 1 MB. The striping granularity can be set for each file to match the properties of its data.

Striping granularity

COARSE: Striping in allocation units (Storage is secured and laid out in allocation units.)

FINE : Striping in finer units (Storage is secured in allocation units and laid out in units of 128 KB.)




Figure 10. Striping Granularity

(5) Mirroring

ASM mirrors at the file level, which can improve high availability. Actually, ASM files are partitioned in allocation units (AU) of 1 MB and these are laid out on different disks to implement mirroring. If a disk fault occurs, file access can be continued without impact by accessing the mirror data on other disks in the disk group.

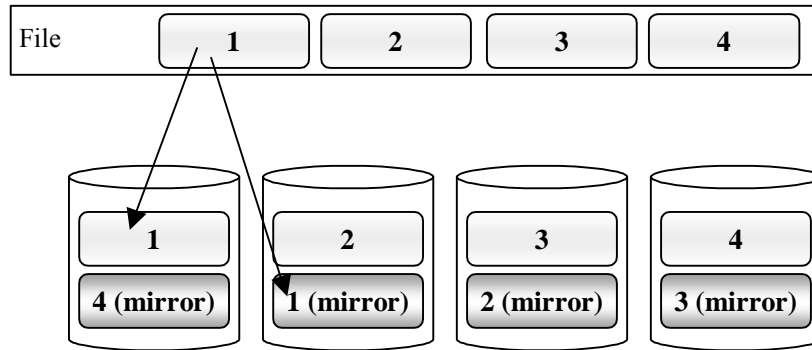


Figure 11. ASM Mirroring Function

Mirroring level

ASM provides two mirroring levels for disk groups.

- Normal (normal redundancy): Data is duplicated.
- High (High redundancy): Data is triplicated.

ASM also provides the ability to create disk groups without mirroring, using external redundancy. This feature allows the customer to leverage the storage's RAID implementation.



Oracle ASM on HDP

Oracle data files include various types of files. The user table space contains a file that will grow over the time correspond to the data growth. The following benefits apply to the virtual volumes of HDP that house the disks which comprise the data files of the user table space.

■ Improved initial costs of storage ownership and capacity design.

All we need is preparing pool volumes with amount initially required. Then by setting virtual volumes with amount expected in the future, creating ASM disk groups with virtual volumes as ASM disks and creating datafiles with AutoExtend on setting, all we have to do is monitoring pool volumes usage for capacity management.

These facts make it possible to say that HDP and ASM can reduce initial introduction costs and lighten the load of capacity design.

Investigate the maximum capacity expected for the Oracle database table space, set this to the virtual volume by means of HDP, and install the number of physical disks to the pool volume on which the data is actually stored. This HDP virtual volume is assigned to the Oracle ASM disk group and the tablespace with the initial minimum required capacity is created in this disk group with the AutoExtend^{*1} set to ON. In this way, as the amount of data increases and the free space in the tablespace becomes inadequate, more disk area is automatically assigned from the remaining capacity in the disk group, extending the tablespace.

These facts make it possible to say that HDP and Oracle database ASM linking can reduce initial introduction costs and lighten the load of capacity design.

■ Improvement of disk management for online disk addition.

With the traditional technology, when the currently allocated Oracle data file space runs out, administrators need to set up on the OS after adding new physical disks, and then ASM rebalance processing starts automatically when ASM rebalance function is set. But by using the HDP and Oracle ASM, you can add the capacity for Oracle data file without setting up on the OS and ASM rebalance process.

In case that virtual volume is set as large capacity than actual HDP pool volume, and when HDP pool volume runs out the capacity, storage will automatically assign physical disk capacity as needed non disruptively since virtual volumes and the available physical disk space are constantly monitored.

Conventionally (using ASM)
In case that current space are runs out on a pool volume.

[Step 1] Addition of HDDs to storage sub system



[Step 2] disk recognition by OS



[Step 3] Chang of permissions and ownership



[Step 4] disk addition to disk group& ASM rebalancing

Using ASM on HDP
· **No Need for ASM rebalances.**
· **No Need for disk recognition by OS.**

Using ASM on HDP
In case that current space are runs out on a pool volume.

[Step 1] Addition of the HDDs to storage sub system.



[Step 2] Addition of LDEVs to pool. volumes



■ Improved capacity management cost

Administrators only need to monitor total capacity consumed by the Oracle database and actual consumed storage capacity instead of monitoring each capacity of Oracle table spaces. This can reduce the capacity management cost.

■ Improved performance design cost

With traditional technology, to realize high performance Oracle system, characteristics of each table space of Oracle database and physical configuration in storage sub-system need to be taken into account during the performance planning. But with using the Oracle ASM on HDP, high performance Oracle system can be easily deployed only with pool volume design since HDP automatically spreads the I/O load of Oracle database accessing the common pool across the available spindles.

*1 AutoExtend

AutoExtend is the automatic extend function of table spaces. When the AutoExtend is applied, in case that table space and table runs out as the data is inserted, Oracle automatically extends the table space and avoids the error caused by short of capacity. When using Oracle ASM and HDP, in case that the table space and table runs out, capacity size required to the table space are allocated to HDP pool volume automatically. The feature will improve the management cost on the capacity expansion.

Best practice of Oracle ASM with HDP

This chapter explains best practice of Oracle and Hitachi storage configuration and backup operation for using Oracle ASM and HDP. Figure 12 and Figure 13 are examples of the best practice.

■ Approach to a system configuration design

[1st Step] Database capacity design

First, design the database capacities.

- (1) Design the maximum Oracle database size with increasing the capacities in future.
- (2) Design the necessary size when setting Oracle data files initially.
- (3) Design the capacities of the files (in which the virtual volumes are not used, such as REDO log files, control files, archive files and so on) except the data files.

Figure 12 and Figure 13 show the examples of the Configuration Diagrams with assuming that the maximum database file size should be about 8TB and the data file size when installing should be about 400GB.

[2nd Step] ASM design

When using ASM on HDP and backup the Oracle database using the Hitachi ShadowImage software (ShadowImage) which is a disk mirroring feature of USP V. it is recommended to create at least three ASM disk groups, as in the table below.

Table 1. Disk Group Configuration When using ASM on HDP

Disk group name	Storage data
DATADG	Data files (System tablespace, SYSAUX tablespace, user table space, temporary tablespace, undo tablespace, server parameter files, etc.)
REDODG	Redo log files, control files #1
ARCHDG	Archive log files, control files #2

For the disk groups, set the ASM external (external mirroring) option. With the configuration of CHA^{*3}, DKA, and the RAID level, you can reduce the risk of data loss.

^{*2}Control package that controls data transfer between cache memory and disk drives

^{*3}Control package that controls data transfer between channels and cache memory

[3rd Step] HDP virtual volume design

Depending on the future maximum capacity reviewing at [1st Step](1), design the virtual volume. Please note that the maximum capacity is 2TB for each 1 volume. With setting multiple virtual volumes, it is possible to make one group as a certain ASM disk group.

Figure 12 and Figure 13 show that the virtual volumes are set 4 volumes, one is 2TB for each.

[4th Step] HDP pool volume design

- Pool volume capacity design

- Prepare a pool volume (the volume on which the actual data is stored) capacity for table space of Oracle database. The capacity needs to be larger than initial requirement.

- Pool volume configuration design


- Design the HDD types, RAID levels, and number of RAID groups from the required pool volume capacity and, confirm if the required performance can be achieved with the hardware spec and configuration.

- If necessary, consider adding RAID groups.

- It is recommended to configure HDP pool with multiple RAID groups consist from same RAID level and HDD type. When multiple DKAs ^{*2} are available, configure pool volumes from the RAID groups under the each DKAs to distribute the load.

- It is recommended to allocate the pool volumes to a RAID group that is different from regular volume (i.e. do not allocate HDP pool volume and non-HDP pool volume in a RAID group).

Figure 12 and Figure 13 show that the pool volumes are set about 1.6TB which is about four times of the data files when setting up initially.



Number of CHAs (Channel Adapters) : 2

Number of DKA (Disk Adapters) : 4

HDD Type : 146GB HDD type

RAID level: RAID5 (3D+1P)

Number of RAID groups (parity groups (PG)): 16

Execution capacity for one RAID group: about 400 GB

*4 Only required for cold backup

Figure 12 Striping USP V Physical Configuration Summary Diagram




Figure 13. Logical Configuration Diagram

■ Performance

Figure 13 (a configuration of HDP and Oracle ASM) and Figure 14 (a typical ASM configuration that has been deployed) show system configurations for performance comparison. Performance test for Oracle system is conducted based on the TPC-C with assuming the regular use case (i.e several MB/s of REDO log be generated). For pool volumes, the recommended configuration of Figure 12 and of Figure 13 comprises four RAID groups, but performance was also measured for a configuration of one RAID group for comparison.

Figure 15 shows these results with the transaction performance value [tpmC] and response time with typical configuration as 100. The performance difference between the typical Oracle ASM configuration and the recommended configuration of Oracle ASM with HDP is a few percent. For normal use there is no issue. But when the pool volume is made up from one RAID group, the transaction performance and response time overhead is about 20-25% compared to a typical ASM configuration, so it is necessary to take into account for the system design.

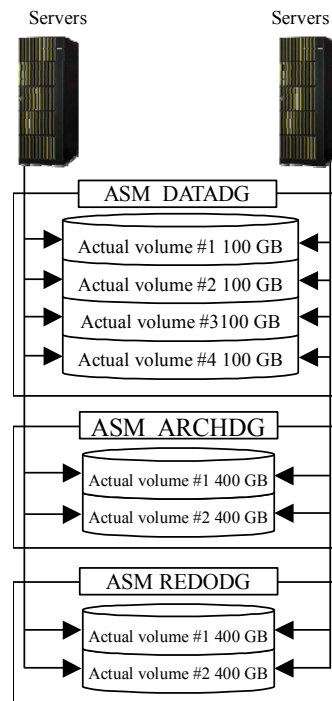


Figure 14. Logical Configuration without HDP (Conventional Type)




Figure 15. Results of Performance Comparison between with and without HDP

■ Backup & restore

Even when the storage is configured linking HDP and Oracle database ASM, the same as for normal Hitachi Storage volumes, HDP virtual volumes can be backed up using Hitachi Storage ShadowImage, which is the Hitachi Storage in-chassis disk mirroring function. In this case, it is recommended that the backup destination volume and the backup source volume have the same configuration and that HDP virtual volume configuration be used. (See the example in Figure 13.)

Below are shown the flows and procedures for cold backups and hot backups, respectively, carried out at specified dates and times as well as one example of the recovery procedure and flow.

Before starting backup operation using ShadowImage, it is necessary to install RAID Manager, set the configuration definition file, etc., then create the pairing for the backed up volume stored on the USP V. For cold backup operation, pairs are created for all volumes belonging to all three disk groups. For hot backup, pairs are created for disk groups storing archive log files and for HDP virtual volumes belonging to disk groups storing data files. Below is an example of command execution. Appendix C shows an example of hot backup of configuration definition files.


- Pair creation execution command example for before cold backup operation

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

- Pair creation execution command example for before hot backup operation

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

(1) Cold backup



*5 After a ShadowImage pair control command (pairsplit, pairresync, etc.) is executed, it is necessary to monitor whether or not the status transition of the ShadowImage pair targeted by the command is completed. The appendix gives an example of a specific script for ShadowImage pair status monitoring.

Figure 16. Logical Cold Backup Flow

• Cold backup procedure

[Step 1] Stopping database on all nodes

Stop the database on all nodes. (shutdown normal or transactional or immediate)

```
SQL> shutdown immediate;
```

[Step 2] Stopping the ASM instance on all nodes

Stop the ASM instance on all nodes. (shutdown normal or transactional or immediate)

```
$ export ORACLE_SID=+ASM1
$ sqlplus / as sysdba
SQL> shutdown immediate;
```

[Step 3] Resync of ShadowImage pair targeted for backup

The pair is created beforehand with ShadowImage and this resyncs the pair that has been targeted for cold backup in the suspended state. When the pair is in the suspended state, update transactions not reflected on the secondary volume are reflected.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairresync -g SIALLDG
```


[Step 4] Suspension of ShadowImage pair targeted for backup

After the resyncing of the ShadowImage pair in [Step 3] is complete, that pair is put in the suspended state.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairsplit -g SIALLDG
```


(2) Restoration using cold backup

Shown below is the flow and procedure for recovery from a cold backup assuming that a certain disk group has been corrupted.



*5 After a ShadowImage pair control command (pairsplit, pairresync, etc.) is executed, it is necessary to monitor whether or not the status transition of the ShadowImage pair targeted by the command is completed. The appendix gives an example of a specific script for ShadowImage pair status monitoring.

Figure 17. Restoration Flow using Cold Backup

- Procedure for restoration using cold backup

Shut down database

Stop the database on all nodes. (shutdown normal or transactional or immediate)

```
SQL> shutdown immediate;
```

[Step 2] ASM instance stopping

Stop the ASM instance on all nodes. (shutdown normal or transactional or immediate)

```
$ export ORACLE_SID=+ASM1
$ sqlplus / as sysdba
SQL> shutdown immediate;
```

[Step 3] ShadowImage pair restoration

This restores the cold backup.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairresync -g SIALLDG -restore -fq normal
```

[Step 4] ShadowImage pair suspension

After the restoration in [Step 3] is complete, the ShadowImage pair is put in the suspended state.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairsplit -g SIALLDG
```

[Step 5] Database startup

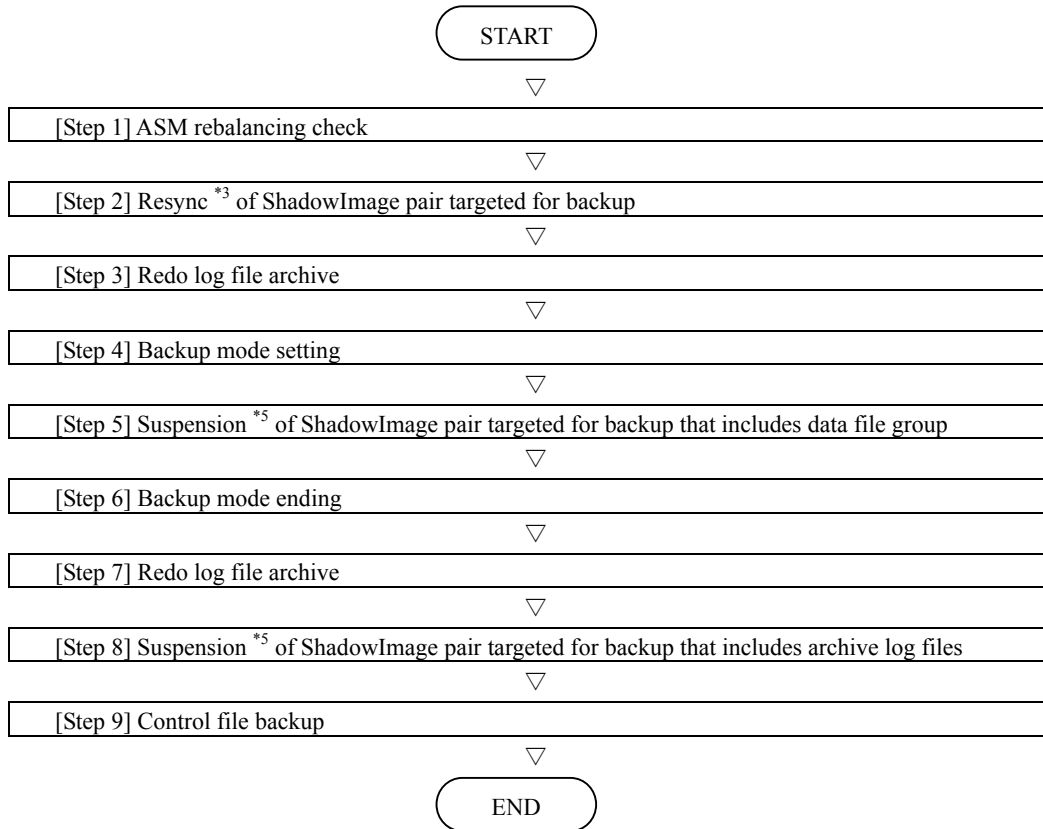
When recovering to the point of the cold backup, database recovery does not occur.

After the suspension of the ShadowImage pair in [Step 4] is complete, start the ASM instance then the database instance in that order.

```
$ export ORACLE_SID=+ASM1
$ sqlplus / as sysdba
SQL> startup

$ export ORACLE_SID=hdpbk1
$ sqlplus / as sysdba
SQL> startup
```

(3) Hot backup



*5 After a ShadowImage pair control command (pairsplit, pairresync, etc.) is executed, it is necessary to monitor whether or not the status transition of the ShadowImage pair targeted by the command is completed. The appendix gives an example of a specific script for ShadowImage pair status monitoring.

Figure 18. Hot Cold Backup Flow

• Hot backup procedure

[Step 1] ASM rebalancing check

Check that ASM rebalancing processing has not been performed. If a hot backup is performed during ASM rebalancing, the ASM meta data can not be matched and an invalid backup may occur. Check the rebalancing processing with the select command below. When rebalancing processing has not been performed, lines such as the following are not selected.

```
$ export ORACLE_SID=+ASM1
$ sqlplus / as sysdba
SQL> select operation,state,group_number,power from v$asm_operation;
no rows selected
```

[Step 2] Resync of ShadowImage pair targeted for backup

The ShadowImage pair that includes a disk group with a data file group and archive log files is resynced in the suspended state.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairresync -g DATADG
# pairresync -g ARCHDG
```

[Step 3] Redo log file archive

After the resyncing of the ShadowImage pair in [Step 2] is complete, the current redo log is archived. This is done to shorten the time spent on the log switch after the end of the backup. It is done before the database is put into backup mode.

```
$ export ORACLE_SID=hdpasm1
$ sqlplus / as sysdba
SQL> alter system archive log current;
```

[Step 4] Backup mode setting

In order to obtain a hot backup, the database is moved to backup mode.

```
$ export ORACLE_SID=hdpasm1
$ sqlplus / as sysdba
SQL> alter database begin backup;
```

[Step 5] Suspension of ShadowImage pair targeted for backup including data file group

The targeted ShadowImage pair that includes a data file group is put in the suspended state.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairsplit -g DATADG
```

[Step 6] Backup mode ending

After the suspension in [Step 5] is complete, the database is put in backup mode.

```
$ export ORACLE_SID=hdpasm1
$ sqlplus / as sysdba
SQL> alter database end backup;
```

[Step 7] Redo log file archiving

The current redo log files for all nodes at this time are archived and output. This secures the log information generated up to the end of the backup. You can check in the alert log for the database instance at all nodes whether or not the archiving was completed correctly.

```
$ export ORACLE_SID=hdpasm1
$ sqlplus / as sysdba
SQL> alter system archive log current;
```

[Step 8] Suspension of ShadowImage pair including archive log files

After the completion of the archiving in [Step 7] has been verified with the alert logs, the pair is suspended in order to backup the archive log area.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairsplit -g ARCHDG
```

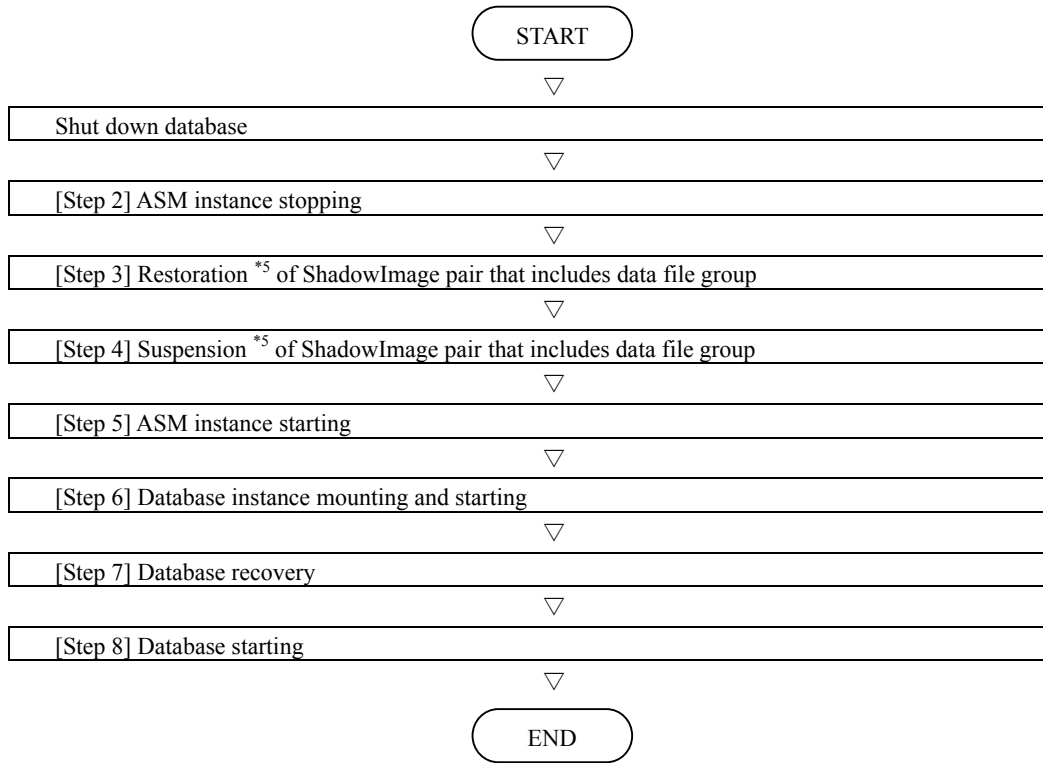
[Step 9] Control file backup

After the suspension of the ShadowImage pair in [Step 8] is complete, the control file backup is acquired separately. Here, we introduce the procedure for acquiring the control file backup in binary format.

```
SQL> alter database backup controlfile to '/home/backup/CTRL.bak';
```

(3) Restoration using hot backup

Below are the flow and procedure for recovery from a hot backup assuming that a disk group that includes data files has been corrupted.



*5 After a ShadowImage pair control command (pairsplit, pairresync, etc.) is executed, it is necessary to monitor whether or not the status transition of the ShadowImage pair targeted by the command is completed. The appendix gives an example of a specific script for ShadowImage pair status monitoring.

Figure 19. Restoration Flow Using Hot Backup

● Procedure for restoration using hot backup

[Step 1] Shut down database

Stop the database on all nodes. (shutdown normal or transactional or immediate)

```
SQL> shutdown immediate;
```

[Step 2] ASM instance stopping

Stop the ASM instance on all nodes. (shutdown normal or transactional or immediate)

```
$ export ORACLE_SID=+ASM1
$ sqlplus / as sysdba
SQL> shutdown immediate;
```

[Step 3] Restoration of ShadowImage pair that includes a data file group

Restore the ShadowImage pair that includes a data file group on which hot backup was performed.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairresync -g DATADG -restore -fq normal
```

[Step 4] Suspension of ShadowImage pair that includes data file group

After the restoration of the ShadowImage pair in [Step 3] is complete, the ShadowImage pair that includes a data file group is put in the suspended state.

```
# HORCC_MRCF=1;export HORCC_MRCF
# HORCMINST=0;export HORCMINST
# pairsplit -g DATADG
```

[Step 5] ASM instance starting

After the suspension of the ShadowImage pair that includes a data file group in [Step 4] is complete, the ASM instance is started.

```
$ export ORACLE_SID=+ASM1
$ sqlplus / as sysdba
SQL> startup
```

[Step 6] Database instance mounting and starting

Start with the database in the mounted state.

```
SQL> startup mount;
```

[Step 7] Database recovery

Execute database recovery processing.

```
SQL> recover database;
```

[Step 8] Database starting

Change the database in the mounted state to the open state.

```
SQL> alter database open;
```



Important observation

Following are the main items to keep in mind when using Oracle ASM on HDP.

(1) Monitoring usage of HDP pool volume capacity

When using Oracle database ASM on HDP, it is necessary to monitor the usage of pool volume capacity so that not to runs out the capacity. On the Hitachi Storage, you can manage threshold values for free pool volume capacity from the viewpoint of the pool volume and from the viewpoint of the virtual volume. When a set threshold value is exceeded, an alert will be reported. Such an alert report can be posted with an SNMP trap. In case that the pool volume capacity runs out, write access to the virtual volume causes an error. Data that has already been written to the pool volume can be accessed and read without any problem.

(2) When a usage effect is not forecast

When HDP is applied to the Oracle files which capacity will not grow(ex: capacity is constant and stable or vary within some capacity), you cannot get the benefit as described in chapter 4.

(3) Areas in HDP pool volumes

When a HDP pool volume is expanded because of the write access to the virtual volume, even if you delete the data, the HDP will not be returned the capacity to the pool volume automatically. To release it operation to the Hitachi Storage is required.

(4) Addition of HDP pool volume

When adding the physical disk because of the short of pool volume, it is recommended to use the same RAID group as the existing HDP pool volume is configured for the performance.

(5) Disk failure case for virtual volume of HDP.

In case that a virtual volume of HDP is used and its disk failure occurred, the same failure phenomena occurs as normal volume.

In running out of HDP pool volume capacity, when Oracle database try to allocate the new resource from pool volume, they will get I/O error and this transaction will rollback.

Appendix A: USP-V Specifications

Table 2. USP V Specifications Table

Item		Specifications	
Product name		Hitachi Universal Storage Platform V	
Maximum number of disk drives that can be mounted		512	
Disk drive type		FC	SATA
Disk drive capacities supported		300 [GB]/146 [GB]/73 [GB]	750 [GB]
Maximum capacity ^{*6}	Sub-system internal storage	332 [TB]	850 [TB]
	Sub-system external storage	247 [PB]	
Cache memory maximum capacity ^{*7}		256 [GB]	
Shadow memory maximum capacity ^{*7}		32 [GB]	
Maximum number of channels connected	Fiber channels	224	
	Main frame fiber channels	112	
	Main frame serial channels	112	
Maximum number of logic paths for main frame connection	Per port	Main frame serial channels: 128 Main frame fiber channels: 65,536	
	Per sub-system	Main frame serial channels: 12,288 Main frame fiber channels: 131,072	
Data path bandwidth		106.4 [GB/s]	
RAID level		RAID5 (3D+1P, 7D+1P), RAID1 (2D+2D, 4D+4D), RAID6(6D+2P)	
Power supply input		Three phase/single-phase 200 V	
External dimensions (W×D×H) [mm]		782 - 3,382×925×1,920	
Noise [dB] ^{*8}		62 dB max.	
Labeled according to the Law Concerning the Rational Use of Energy (FY2007 regulations)	Section	i	i
	Energy consumption efficiency ^{*9}	0.095	0.034

^{*6} This capacity is the value calculated with 1 TB = 1,000⁴ bytes.

^{*7} This capacity is the value calculated with 1 GB = 1,024³ bytes.

^{*8} This is the value measured at a position 1 meter from the door on the floor of the sub-system.

^{*9} The energy consumption efficiency is the power consumption measured using the measurement method prescribed by the Law Concerning the Rational Use of Energy multiplied by the memory capacity prescribed by the Law Concerning the Rational Use of Energy.

Appendix B: HDP Support Configuration Summary

Table 3. HDP Support Configuration Summary

#	Item	Support configuration
1	Number of virtual volumes	4096 volumes/pool
2	Virtual volume size	46 MB - 2.0 TB
3	Maximum number of pools	32
4	Number of pool volumes	256/pool
5	Pool volume size	8 [GB] - 2.0 [TB]
6	Pool volume type	Internal volume, external volume Note: Internal and external volumes can not be mixed within a single pool.

Appendix C: RAID Manager Configuration Definition File Example

■ Example of RAID Manager configuration definition file for hot backup configuration

<pre> horcm0. conf HORCM_MON #ip_address service poll(10ms) timeout(10ms) bk_serv1 40000 12000 3000 HORCM_CMD #dev_name dev_name dev_name /dev/rdisk/c1t0d4s2 /dev/rdisk/c2t1d10s2 HORCM_DEV #dev_group dev_name port# TargetID LU# MU# DATADG dev1 CL1-E 0 0 0 DATADG dev2 CL1-E 0 1 0 DATADG dev3 CL1-E 0 2 0 DATADG dev4 CL1-E 0 3 0 ARCHDG dev5 CL2-F 1 0 0 ARCHDG dev6 CL2-F 1 1 0 HORCM_INST #dev_group ip_address service DATADG bk_serv1 40001 ARCHDG bk_serv1 40001 </pre>	<pre> horcm1. conf HORCM_MON #ip_address service poll(10ms) timeout(10ms) bk_serv1 40001 12000 3000 HORCM_CMD #dev_name dev_name dev_name /dev/rdisk/c1t0d4s2 /dev/rdisk/c2t1d10s2 HORCM_DEV #dev_group dev_name port# TargetID LU# MU# DATADG dev1 CL1-E 0 16 0 DATADG dev2 CL1-E 0 17 0 DATADG dev3 CL1-E 0 18 0 DATADG dev4 CL1-E 0 19 0 ARCHDG dev5 CL2-F 1 16 0 ARCHDG dev6 CL2-F 1 17 0 HORCM_INST #dev_group ip_address service DATADG bk_serv1 40000 ARCHDG bk_serv1 40000 </pre>
---	---

Appendix D: ShadowImage Pair Status Monitoring

After a ShadowImage pair control command (pairsplit, pairresync, etc.) is executed, it is necessary to monitor whether or not the status transition of the ShadowImage pair targeted by the command is completed. The status of a ShadowImage pair is monitored with the value returned by the pairvolchk command. In this way, after a pair control command such as pair resync or pair suspend is executed, you can easily monitor whether or not the pair status transition has completed normally. Table 4 shows typical return values for each ShadowImage pair status.

Table 4. Values Returned by the Pairvolchk Command

#	Pair Status	Pair status description	Primary volume (PVOL)		Secondary volume (SVOL)	
			Read/Write enabled	pairvolchk -ss Return value	Read/Write enabled	pairvolchk -ss Return value
1	SMPL	Volume that has no pair relationship	Read/write enabled	11	Read/write enabled	11
2	PAIR	Duplication status of a volume with a pair relationship.	Read/write enabled	23	Read enabled ^{*8}	33
3	COPY	Unfinished copy from primary volume to secondary volume	Read/write enabled	22	Read enabled ^{*8}	32
4	RCPY	Unfinished restore option copy from the secondary to the primary volume	Read/write enabled	22	Read enabled ^{*8}	32
5	PSUS	The pair status is maintained, but updates to the secondary volume are suspended. Update information is maintained in a differential bitmap.	Read/write enabled	24	Read enabled ^{*8}	34
6	PSUE	In suspended status due to a fault or forcible interruption of a copy.	Read/write enabled*	25	Read/Write enabled ^{*8}	35

^{*8}Depending on options used in the command executed immediately before transition to this state and on whether or not a fault is present, not only writing, but reading may also be disabled.

[Example] Example of simple script for ShadowImage pair monitoring

Example of monitoring script with a 10-second interval until dev_group=DATADG(-g DATADG) becomes a pair, depending on the pairresync command

- Normal end: The pair resync is finished. (Value returned by pairvolchk -ss: 23)

Value returned by script "0"

⇒You can proceed to the next step.

- Abnormal end: The pair resync is unfinished. (Value returned by pairvolchk -ss: other than 22 and 23)

Value returned by script "99"

⇒Investigation of the cause is required.

* If the value returned by pairvolchk -ss is 22, it is judged that the pair resync is underway and the pair status is checked again after 10 seconds.

```
#!/bin/sh
RTN=$?

while [ $RTN -ne 23 ]
do
    sleep 10
    HORCC_MRCF=1;export HORCC_MRCF
    HORCMINST=0;export HORCMINST
    pairvolchk -g DATADG -ss -nomsg
    RTN=$?

    case ${RTN} in
        23)
            echo "pairresync done."
            ;;
        22)
            echo "COPY"
            ;;
        *)
            echo "pairresync error."
            exit 99
            ;;
    esac
done
```