

An introduction to Automatic Storage Management

Christian Bilien



Automatic Storage Management

- ◆ A new feature introduced in Oracle Database 10g
- ◆ Provides a vertical integration of the file system and volume manager for Oracle database files
- ◆ Spreads database files across all available storage for optimal performance and resource utilization
- ◆ Enables simple and non-intrusive resource allocation and provides automatic rebalancing

ASM Installation Best Practices

- ✦ Install ASM in a separate ORACLE_HOME from the database ORACLE_HOME
- ✦ Provides higher availability and manageability
- ✦ Allows independent upgrades of the database and ASM.
- ✦ De-installation of database software can be performed without impacting the ASM instance

ASM Instance

- ◆ Only one ASM instance per node
- ◆ Manages storage for all database instances on the node
- ◆ In RAC environments, each clustered node will have an ASM instance
- ◆ ASM instances use the interconnect to synchronize metadata

Sample init.ora parameters

```
instance_type=asm
#SGA sizing -----
large_pool_size=20M
shared_pool_size=80M
processes=55
#Storage -----
asm_diskstring='ORCL:*'
asm_diskgroups='DATA','FLASH','DATA_SOX'
#RAC parameters -----
+ASM1.instance_number=1
+ASM2.instance_number=2
cluster_database=true
# trace files destinations -----
background_dump_dest=/opt/oracle/product/server32/10gr1/admin/
+ASM/bdump
core_dump_dest=/opt/oracle/product/server32/10gr1/admin/+ASM/
cdump
user_dump_dest=/opt/oracle/product/server32/10gr1/admin/+ASM/_5
udump
```

Setting init.ora parameter

- ✦ Use default values
- ✦ Only processes parameter may need modification
- ✦ Processes parameter setting is dependent on the number of databases that connect to ASM
 $\text{processes} = 25 + 15n$, where $n = \#$ databases connected to ASM

ASM administration

- ◆ Disks
- ◆ Diskgroups
- ◆ Files



Diskgroup Best Practices

- ✦ Create two diskgroup
 - One for database area
 - Second one for flash recovery area
- ✦ Create diskgroups using large number of similar type disks
 - same size characteristics
 - same performance characteristics

Diskgroup Best Practices

- ✦ The size of the flash recovery area diskgroup will depend on what is stored and how much is retained
- ✦ If possible, physically separate the database area and flashback area
- ✦ make sure the two areas do not share the same physical spindles

Diskgroup Best Practices

- ◆ To minimize search overhead, perform all required mount operations in a single mount command
- ◆ If adding or removing multiple disks, make the change in a single rebalance operation.
- ◆ This coalesces rebalance operations and reduces overhead

Diskgroup Best Practices

- ◆ Use ASM External Redundancy when creating disk groups with high end storage arrays
- ◆ Use ASM Redundancy for low-end (modular) or JBOD storage array systems
- ◆ Use failure groups with ASM Redundancy

Disk Best Practices

- ✦ Make sure disks span multiple backend disk adapters
- ✦ Implement multiple access paths to the storage array using two or more HBAs or initiator
- ✦ Deploy multi-pathing software over these multiple HBAs to provide I/O load-balancing and failover capabilities.

Disk Best Practices

- ✦ If using hardware RAID, make LUN stripe size as close to 1mb as possible
- ✦ Use OS disk labels when available
 - Prevents accidental user overwrites of disks
 - Easier management of disks
 - Make sure the disk (partition) starts at 1Mb a boundary, to insure proper I/O alignment

Database-ASM Best Practice

- ✦ Create CRS dependency between database instance and ASM instance
- ✦ Ensures ASM instance starts up before starting database instance
- ✦ Use 10g Automatic Memory Management
- ✦ Allows ASM instance to cleanly shutdown database instance

```
srvctl modify instance -d crm -i crm1 -s asm1
```

Database-ASM Best Practices

- ✦ The following recommendations for database SGA sizing can be used to calculate the
 - ✦ SGA_TARGET value
 - large_pool = Add additional 600k
 - Processes = Add 16

Database-ASM Best Practices

Shared_pool – Aggregate the values from the following queries to obtain current database storage sizes that is either already on ASM or will be stored in ASM. Then determine the redundancy type that is used (or will be used), and calculate the shared_pool, using the aggregated value as input.

- . select sum(bytes)/(1024*1024*1024) from v\$datafile;
- . select sum(bytes)/(1024*1024*1024) from v\$logfile a, v\$log b where a.group#=b.group#;
- . select sum(bytes)/(1024*1024*1024) from v\$tempfile where status='ONLINE';

For diskgroups using external redundancy = (Every 100Gb of space needs 1Mb of extra shared pool) + 2M

For diskgroups using Normal redundancy: (Every 50Gb of space needs 1Mb of extra shared pool) + 4M.

For diskgroups using High redundancy: (Every 33Gb of space needs 1Mb of extra shared pool) + 6M.

Database-ASM Best Practices

- ✦ RMAN is the recommended method for backup/recovery of database
 - RMAN is completely ASM diskgroup aware
 - DataPump dumpsets can be created in diskgroups

ASM Command Line Utility *fnezi*

\$ asmcmd

ASMCMD> **ls -l**

State	Type	Rebal	Unbal	Name
MOUNTED	EXTERN	N	N	DISKGROUP1/
MOUNTED	EXTERN	N	N	FLASH/ASMCMD

ASMCMD> **cd DISKGROUP1**

ASMCMD> **ls -l**

Type	Redund	Striped	Time	Sys Name
				Y SALES/

ASMCMD> **cd SALES**

ASMCMD> **ls -l**

Type	Redund	Striped	Time	Sys Name
				Y CONTROLFILE/
				Y DATAFILE/
				Y ONLINELOG/
				Y PARAMETERFILE/
				Y TEMPFILE/
				N spfilesales.ora =>

+DISKGROUP1/SALES/PARAMETERFILE/spfile.268.569267897ASMCMD

ASMCMD> **cd DATAFILE**

ASMCMD> **pwd**

+DISKGROUP1/SALES/DATAFILE

Migration to ASM

- ◆ Database must already be 10g database
- ◆ Use RMAN to migrate to ASM from raw or file system
- ◆ RMAN can be used to migrate entire databases or individual files

More information

<http://www.oracle.com/technology/products/database/asm/index.html>