



Building Your Storage Infrastructure with 10g & ASM

**Marcus Wilson – Senior Technical Consultant
Tenure**

Agenda

- **About Tenure**
- **10g – The Grid 'Do more with less'**
- **Understanding the GRID**
- **Automation Components**
- **9i – OFM - Oracle File management**
- **ASM - Automatic Storage Management**
- **The ASM Instance**
- **The Create Database with just ASM**
- **The Create Database with ASM and OFM**
- **Questions**

About Tenure – Our Services

- **Tenure specializes in implementing Oracle's technology solutions, offering a full range of consulting and training.**
- **Our team of design architects, project managers, analysts, DBA's and developers will work to ensure a successful adoption of your new solution.**
 - **Custom Web Development; Oracle Portal / JSP**
 - **Legacy Data Migration to Oracle**
 - **Database Administration / Performance Tuning**
 - **RAC / Cluster Implementations**
 - **Full scale Data Warehouse / BI Projects**
 - **Mentoring and Training**
 - **Custom Business Applications**
 - **Project Management**

The
GRID...

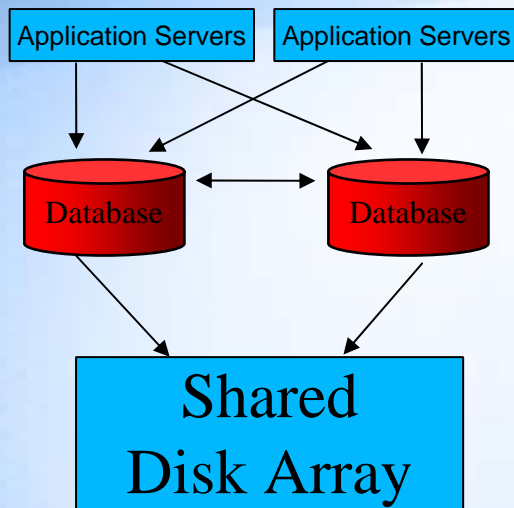
The Hardware
GRID...

The Storage
GRID...

Making the most of your enterprise...

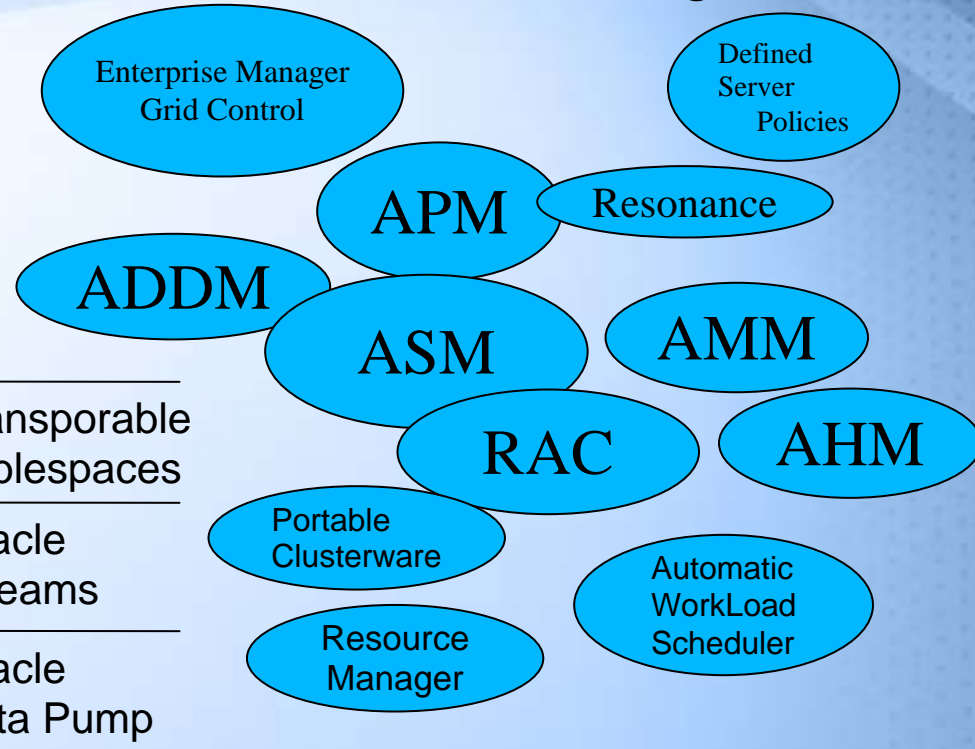
9i comparisons to 10g

Grid in 9i



The overall management components of GRID is what was missing....

Grid in 10g



The 10g Management Components AUTOMATION !

Oracle 10g is by no means just about grid. There are substantial additional features to provide improved database performance, high availability, and easier and more comprehensive management features.

This presentation concentrates on ASM, but the others are worth mentioning:

AWR - Automatic Workload Repository

AMT - The Automatic Maintenance Tasks

ADDM - Automatic Database Diagnostic Monitor

AMM - Automatic Memory Management

STA - SQL Tuning Advisor

SGA - Server Generated Alerts

ASM - **Automatic Storage Management**

AWR - 10g Management Components

AWR - Automatic Workload Repository

The next evolution of STATSPACK... The Automatic Workload repository is the statistics collection agent that supplies data to ADDM - Automatic Database Diagnostic Monitor, the STA -SQL Tuning Advisor, Undo Advisor, and the Segment Advisor.

AMT - 10g Management Components

AMT - The Automatic Maintenance Tasks

The Automatic Maintenance Tasks automate the routine tasks of refreshing statistics as they become stale, and rebuilding indexes when they become sub-optimal.

ADDM - 10g Management Components

ADDM - Automatic Database Diagnostic Monitor

The Automatic Database Diagnostic Monitor (ADDM) taps the performance diagnostics stored in the AWR and provides the DBA with reports that show what problems exist, and many times presenting detailed analysis into the root cause and treatment.

AMM - 10g Management Components

AMM - Automatic Memory Management

Automatic Memory Management is a single parameter (`SGA_TARGET=`) that automates the allocation of RAM between the data buffers, shared pool, and log buffers.

STA - 10g Management Components

STA - SQL Tuning Advisor

The SQL Tuning Advisor (STA) works with the Automatic Tuning Optimizer (ATO) to analyze historical SQL workload (using data from the AWR), and generates recommendations for new indexes and materialized views that will reduce the disk I/O associated with troublesome SQL statements.

SGA - 10g Management Components

SGA - Server Generated Alerts

Server Generated Alerts (SGA) interface with the OS to send e-mail messages when an external problem is impeding Oracle performance.

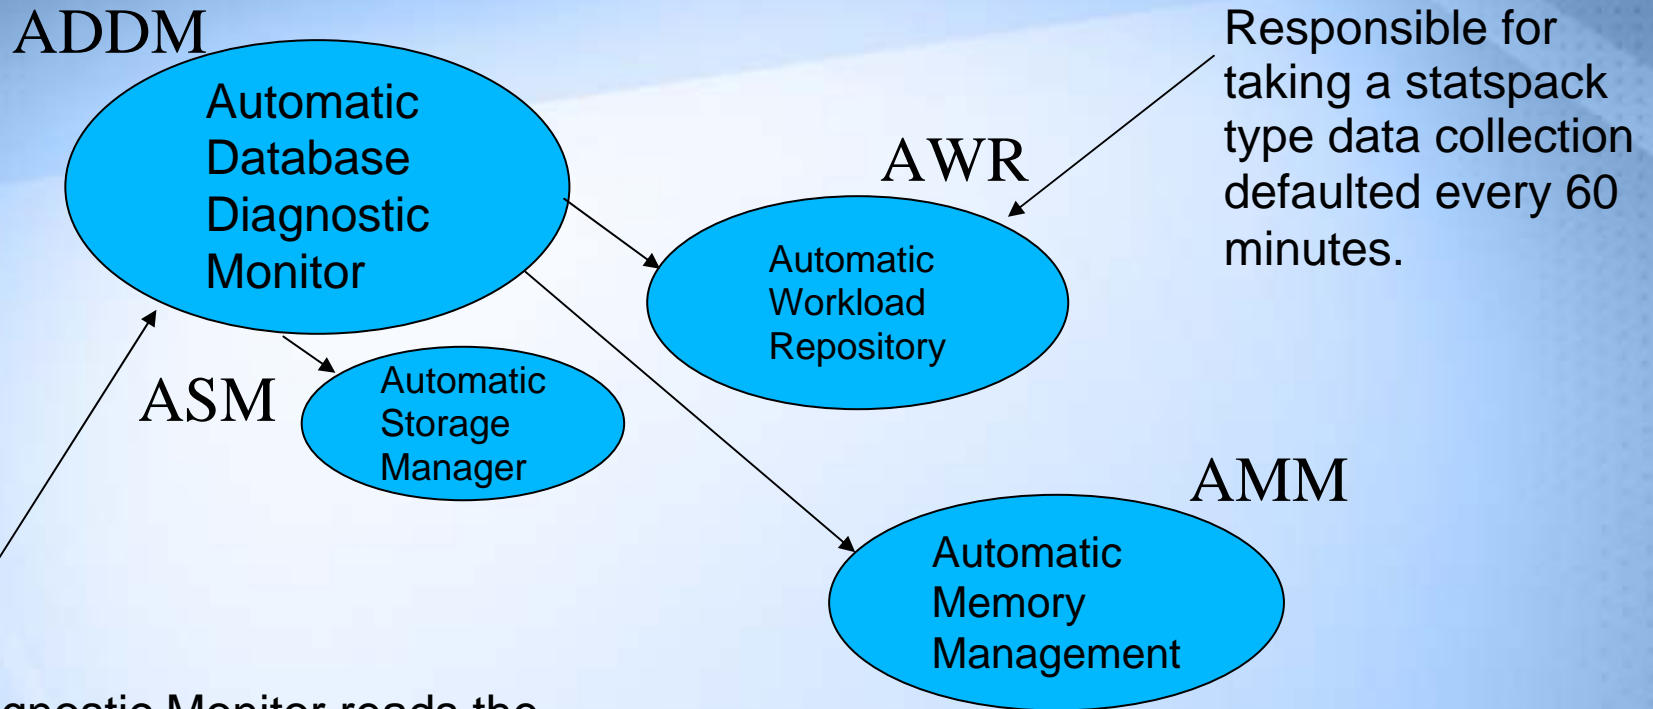
ASM - 10g Management Components

ASM - Automatic Storage Management

ASM will automatically stripe data across the disks in one megabyte chunks. This has the effect of randomly distributing data over a set of disks and hence increasing the average speed of disk reads and writes.

In effect there are no disk volumes! Just a pool of disk storage that is optimized for retrieval. When new disks are added to the pool, ASM automatically re-balances the distribution of data.

Management Components Integrate



Responsible for taking a statspack type data collection defaulted every 60 minutes.

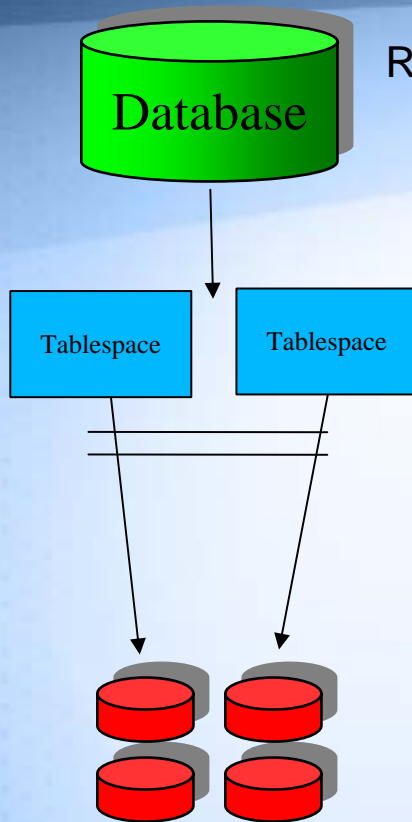
The Diagnostic Monitor reads the data collected by the AWR and uses a built in knowledge base to interact with the DBA and/or Automatic Storage Manager.

It used to be said 'Bigger is better'

SUVs	Have you ever tried to park an Excursion?
Houses	It is way to big after the kids are gone!
Computers	Just what are we going to do with that Cray?
Buildings	Did someone say there was lease space Available?
Airplanes	What zone was that?... Hey I'm in first class!

With the continuing increases in speed and technology, the need for huge multi-processor single box systems is starting to wane. What we have now is the ability to push our smaller enterprise servers to massively parallel computing systems.

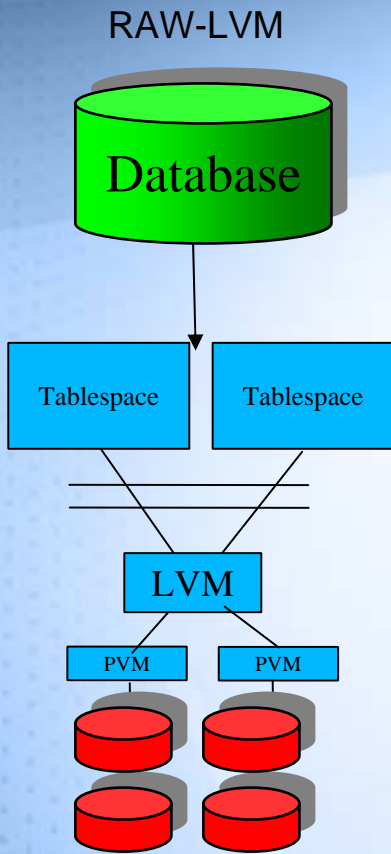
RAW - File Management Options



RAW

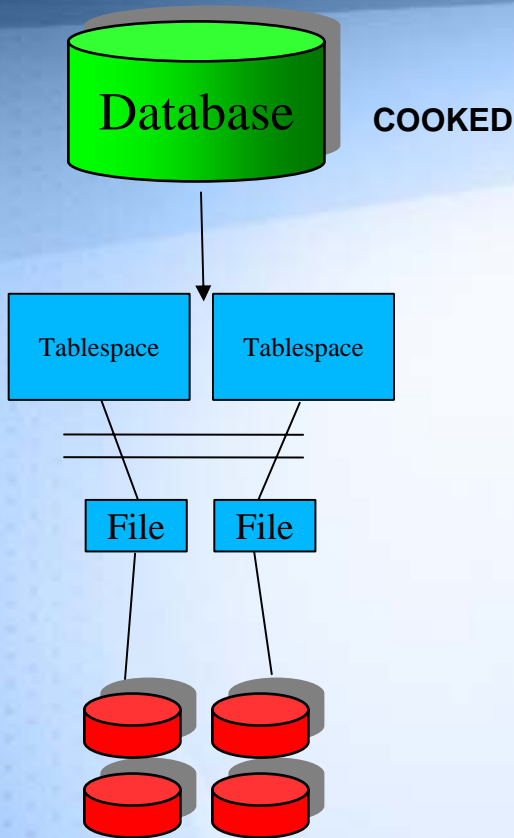
- The DBA and System Administrator spend time allocating the physical raw disk partitions to best utilize the disk. After plans are drawn SysAdmin creates the physical raw partition and presents it to the DBA.
- This process of planning and presenting the raw device is still in use today. RAW devices still show higher throughput than their cooked counterparts.

RAW-LVM - File Management Options



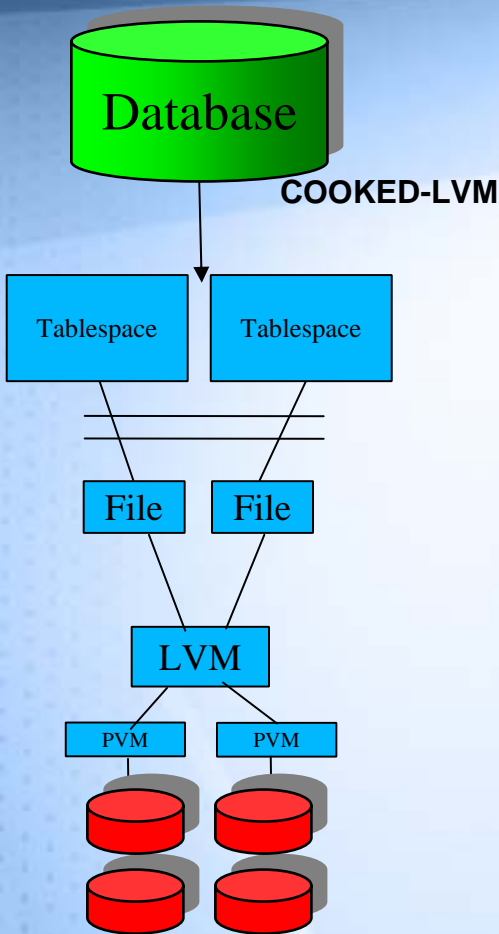
- The introduction of LVM (Logical Volume Management) allowed the introduction of RAID disk management possible beyond simple mirroring.
- The ability to stripe the disks together and present them as a logical presentation to the OS added greater flexibility to the mix.
- However, with the introduction of RAID the concentration on I/O performance and disk contention became the DBAs nightmare.

COOKED - File Management Options



- The cooked file system allowed the System Administrator to present the RAW devices to the Oracle database as mount points. This decreased the complexity of managing the RAW file systems.
- A simple copy command went from the confusing dd command to a simple cp command. Backup scripts, datafile movement scripts were much easier to manage.
- DBA heaven... The system administrator carried the burden of keeping the I/O contention to a minimum through exhaustive planning and execution.

COOKED LVM - File Management Options



When the DBA and System Administrator sit down to carve out a 23TB system, the ability to present a fewer number of LUNS as mount points becomes increasingly popular.

There is a great deal of planning that must be solidified prior to the creation of the LUNS. Once the Logical representations are presented to the OS the installation of the database across the LUNS is possible.

The Logical Volume Manager is fully responsible for management of the physical units.

9i – Introduction of OMF

Question: How many have accepted OMF as a standard?

Some of the issues are:

The original draw back is that we as 'all controlling' DBA's could not manage the I/O contention when we were not able to place the files where we want them.

We mostly used them for small, low transaction databases or development databases. Why, we trusted them for low transactions and it made our life much easier for development.

10g removes the limitations !

For OMF to kick in.....

We have to configure these INIT.ora/SPFILE.ora settings.

- `db_create_file_dest = '/u1/oracle/oradata/testdb'`
- `db_create_online_log_dest = '/u1/oracle/oradata/redo`

Any Limitation?

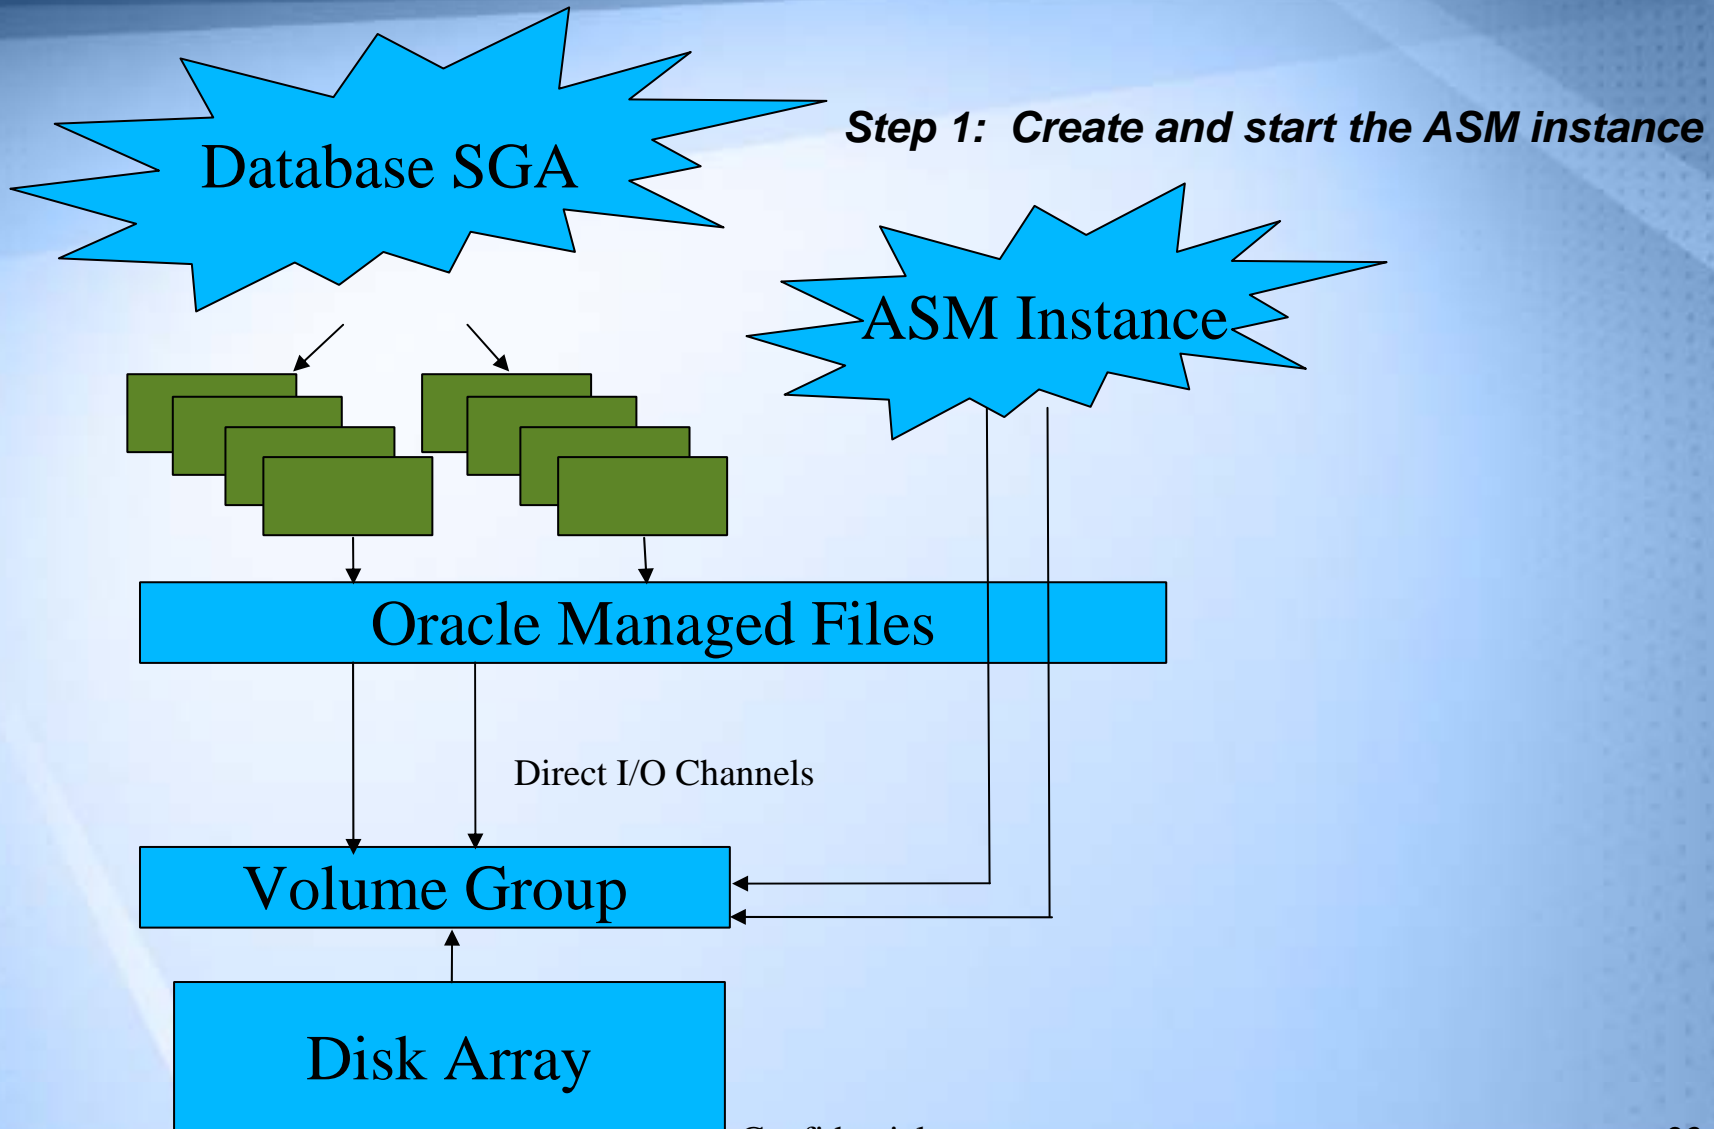
- No RAW devices.
- There is OS level filesystem control.
- The destination can only represent 1 filesystem. However, it can be a single pseudo-disk representing multiple physical disks.

10g Automatic File Management

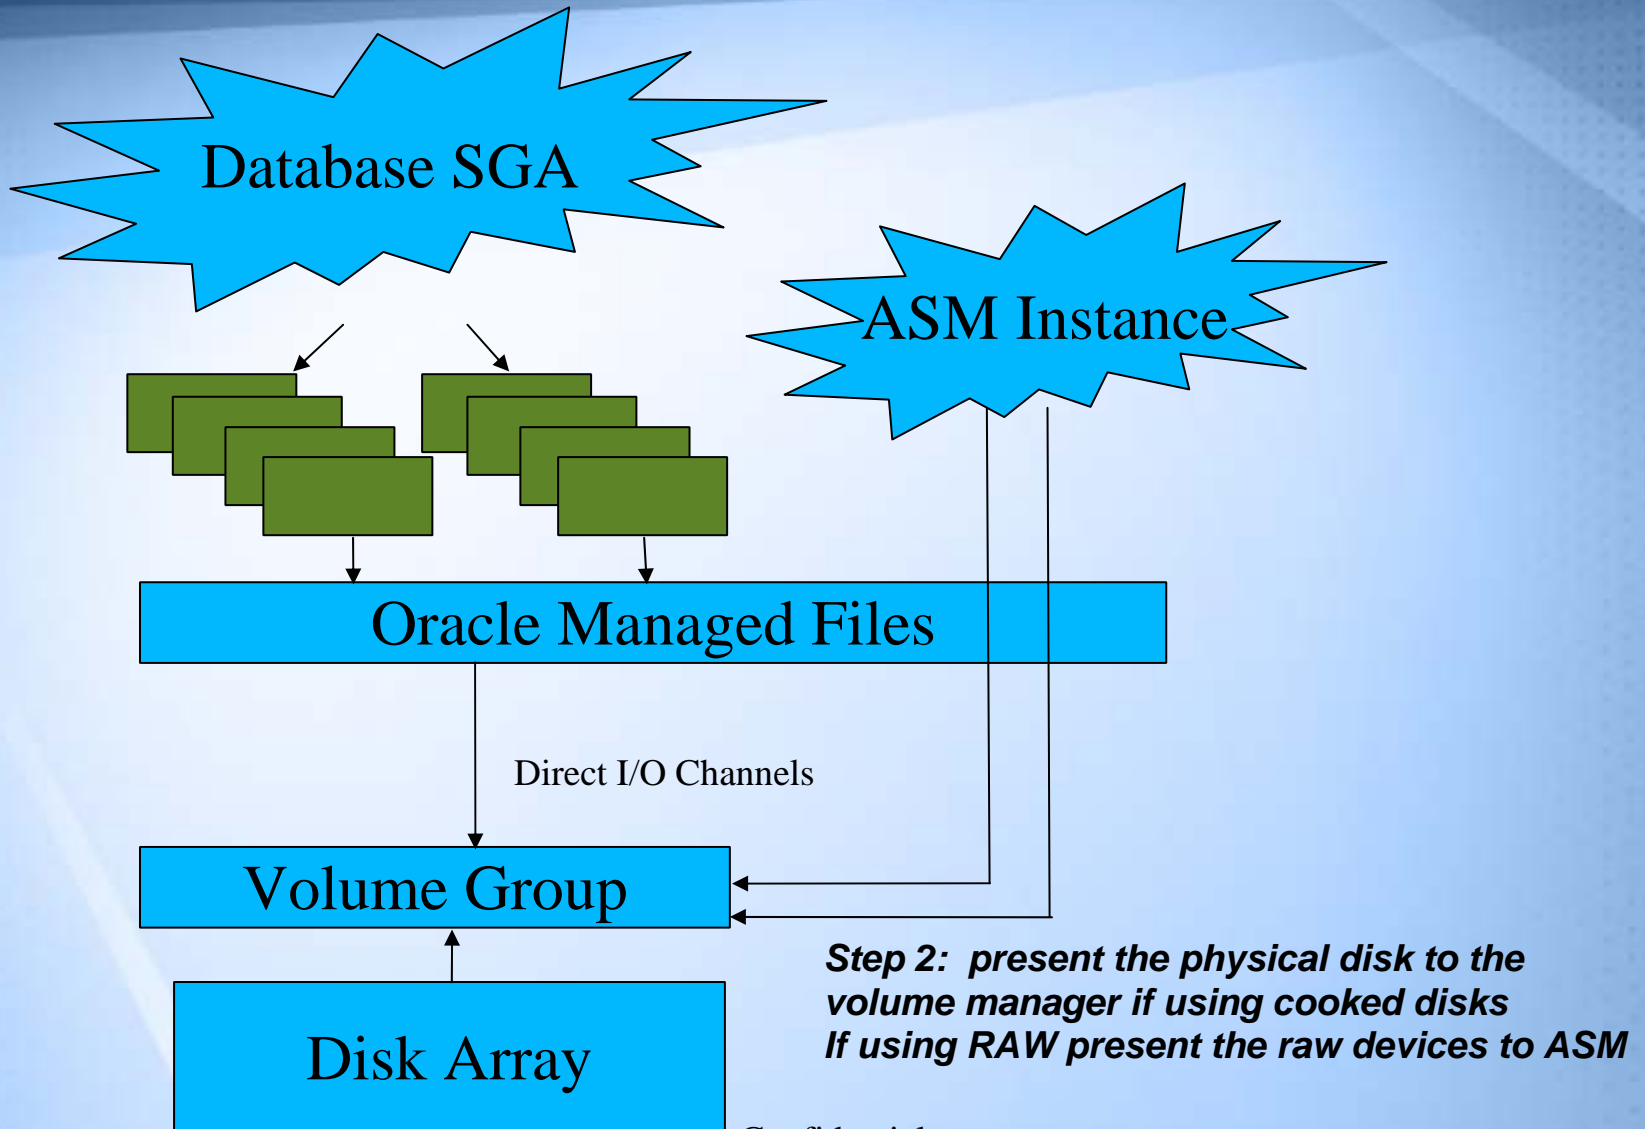
Necessary Components

- The ASM Instance: The instance holds the metadata descriptions of the physical datafile locations and their relation to Disk Groups.
- Disk Groups: The logical grouping of physical devices presented to the ASM Instance.
- A good working relationship between the DBA group and the System Administrators.

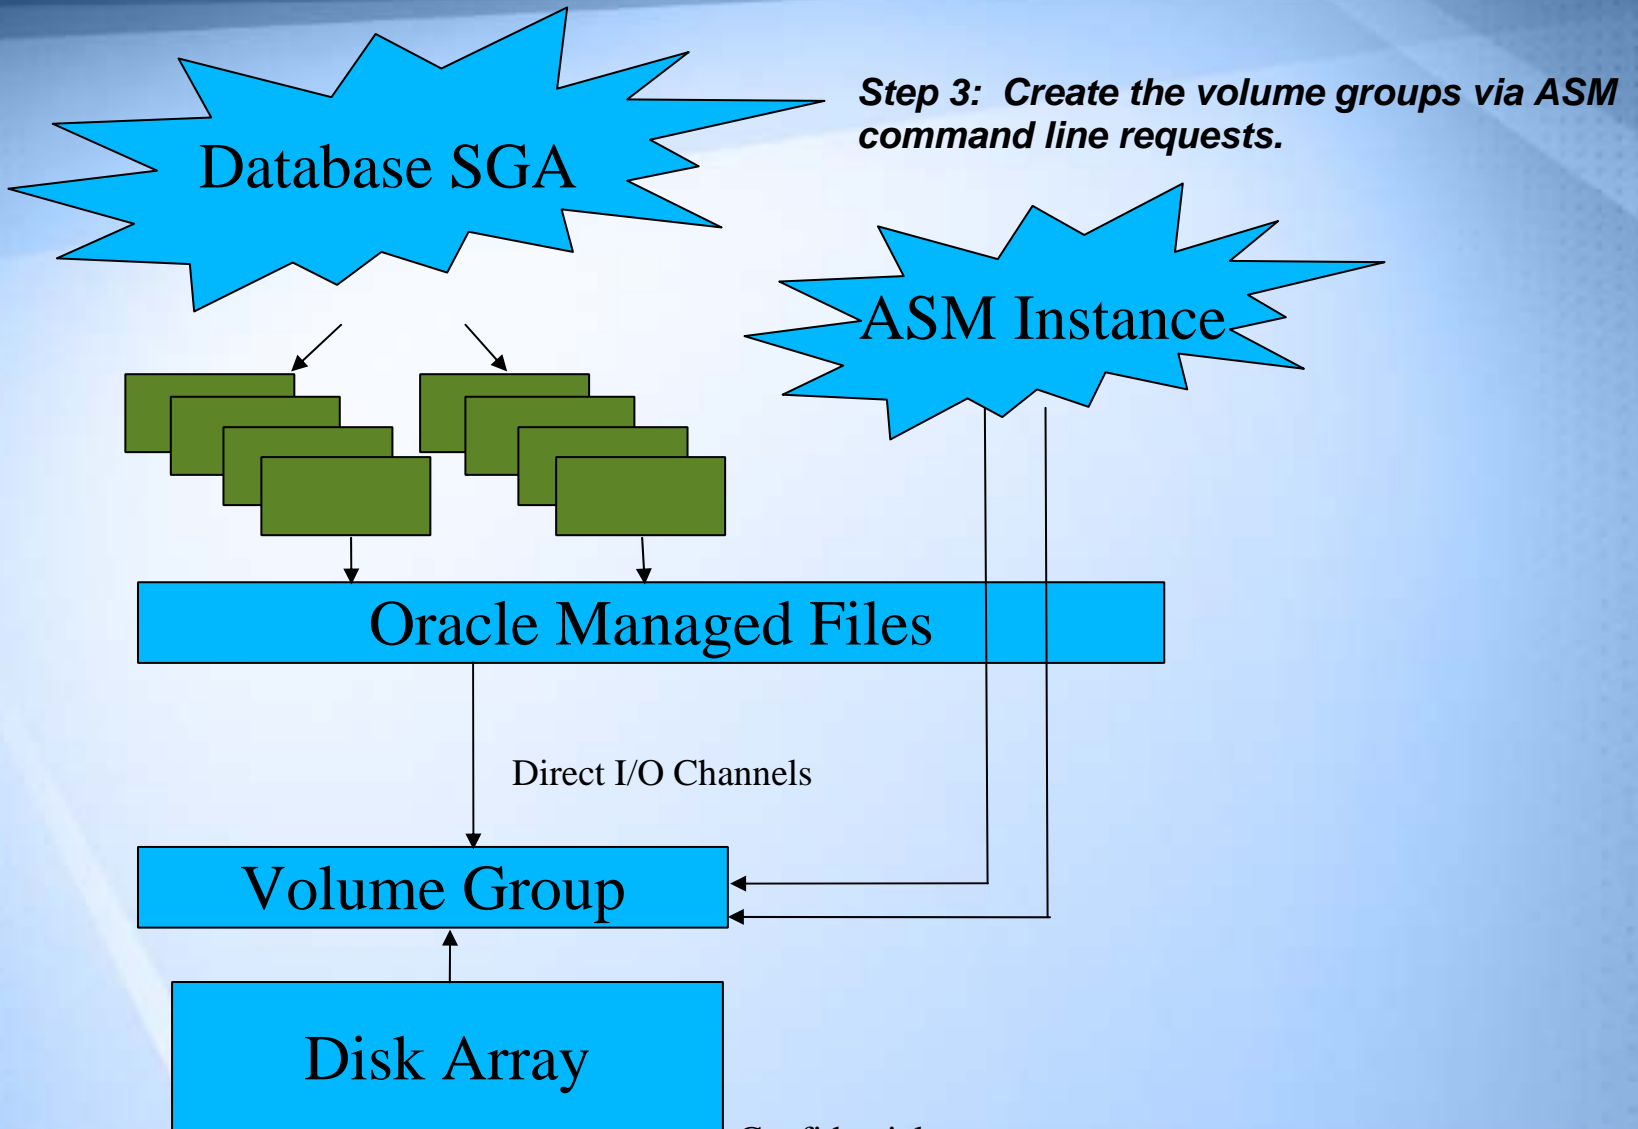
Virtual Tour of ASM



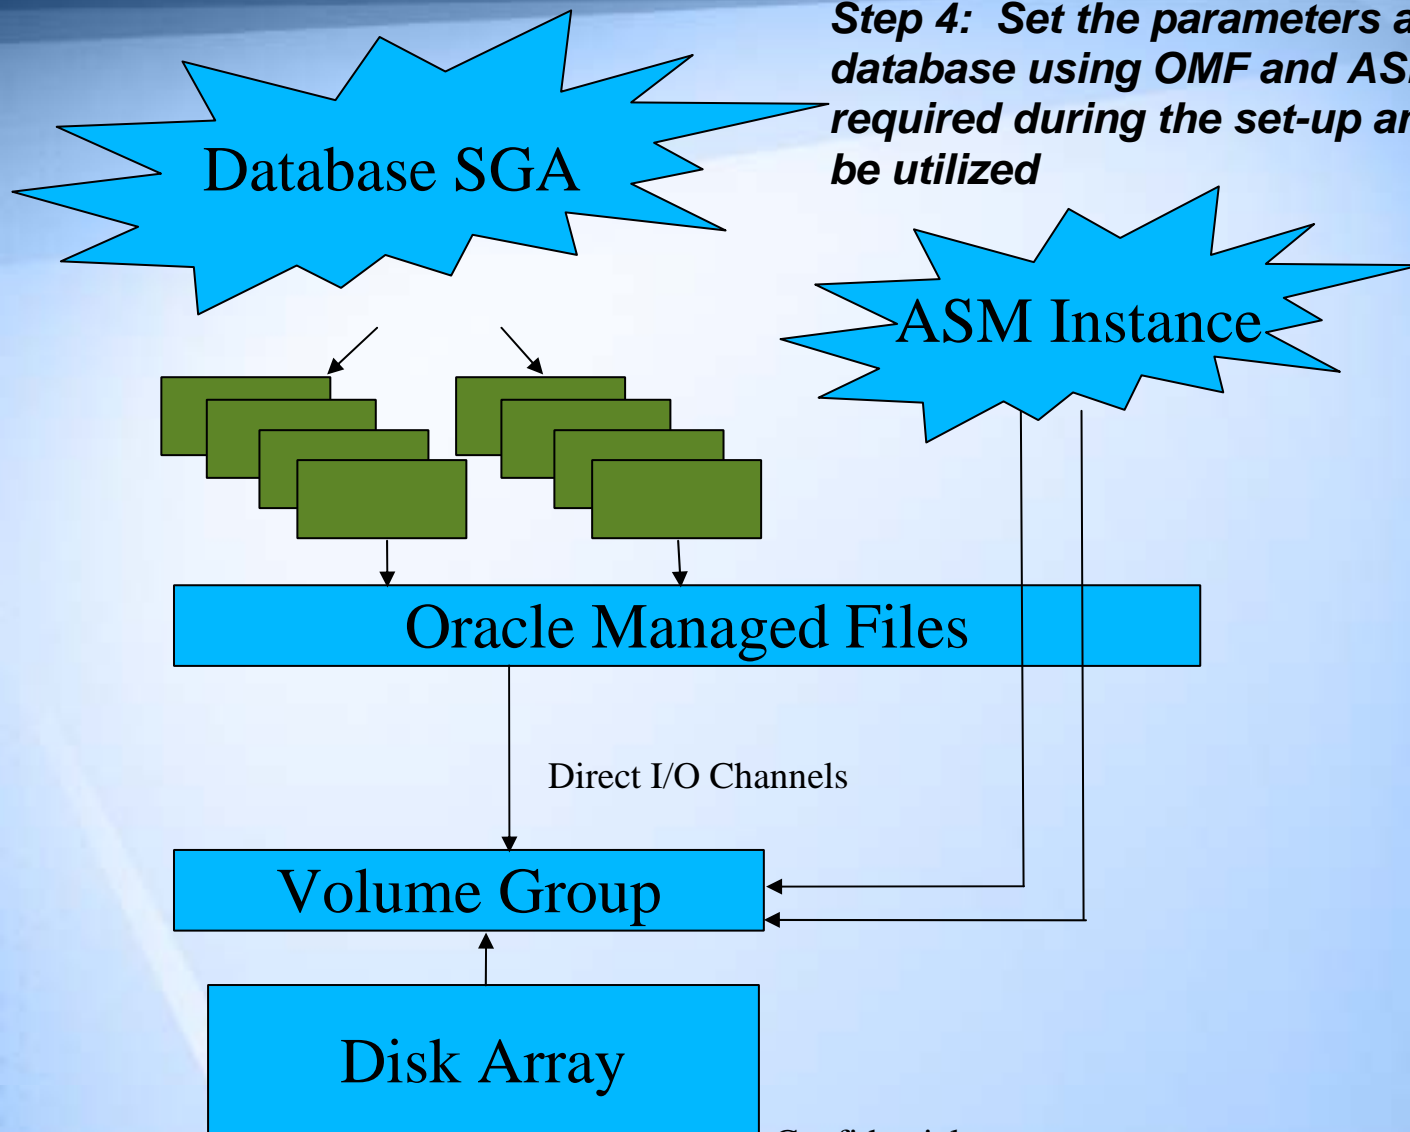
Virtual Tour of ASM



Virtual Tour of ASM

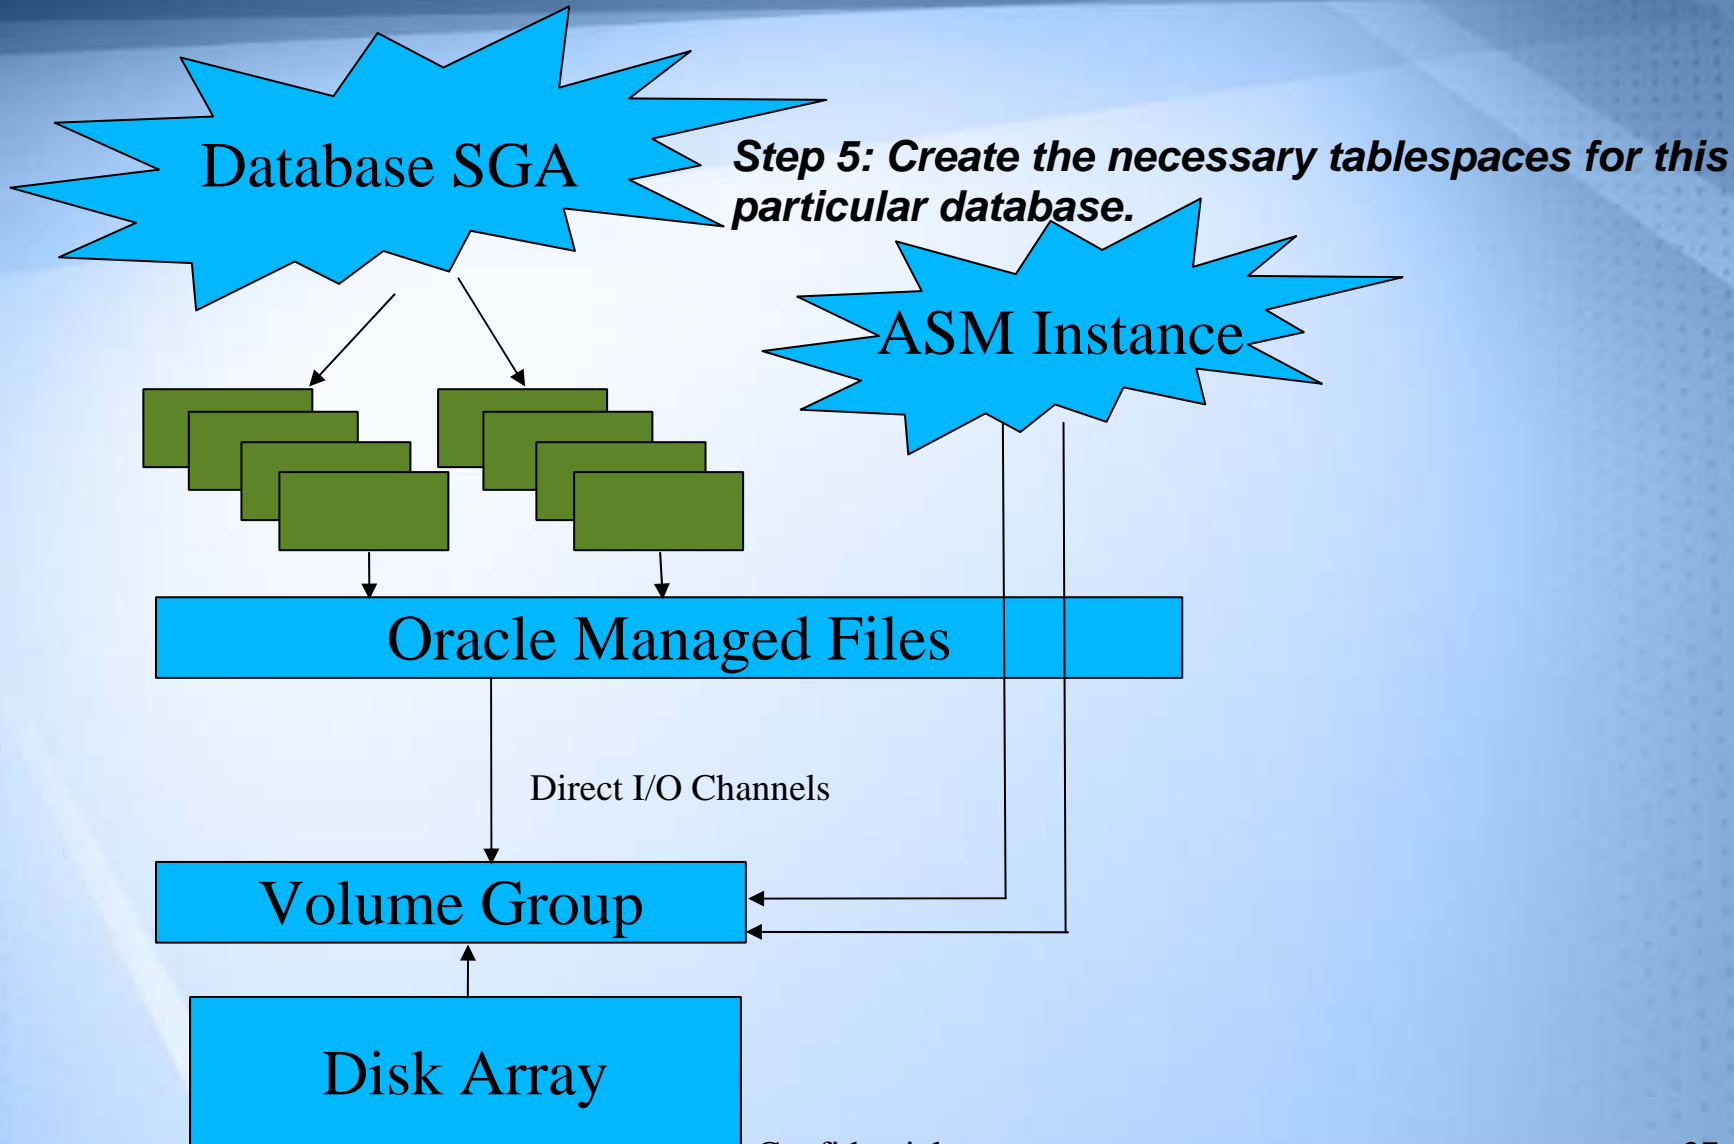


Virtual Tour of ASM



Step 4: Set the parameters and create the database using OMF and ASM. OMF is not required during the set-up and just ASM can be utilized

Virtual Tour of ASM



ASM- Automatic File Management

Highlights to understand....

- Before a database can be opened, the physical files must be located using the controlfile. If the controlfile is part of the disk group, how does the database startup?
- The ASM Instance must be available prior to the database startup procedure and must always be available while the database is up.
- Why you might ask? The ASM mounts the disks and creates an extent mapping, which is then passed on to the database instance.
- The impact of the ASM on performance is minimal. The SGA is a whopping 64MB, unless the system is already overtaxed there is no impact on the memory available to the database instance.

ASM- Automatic File Management

Overview of configuration process

Here are the basic steps in implementing the ASM Instance:

1. Define the ASM instance in oratab:

+ASM:/u1/oracle/product/10.1.0:Y

The “+” sign means that this is an ASM instance. With 10g if you see a “+” sign automaticall think ASM. Be sure to think though any scripts that you have written that access the oratab file for this new feature.

Set up the files known to the operating system.

Remember that you cannot access ASM files directly from the OS. Therefore, the Alert_logs, the core dumps, .trc files must be in an OS filesystem and readable.

ASM- Automatic File Management

Overview of configuration process

(continued)

The pfile must contain:

***.asm_diskstring='/dev/raw/*'**

***.background_dump_dest='/u1/oracle/admin/+ASM/bdump'**

***.core_dump_dest='/u1/oracle/admin/+ASM/cdump'**

***.db_unique_name='+ASM'**

***.instance_type='ASM'**

***.user_dump_dest='/u1/oracle/admin/+ASM/udump'**

ASM- Automatic File Management

Overview of configuration process

(continued)

Start the ASM.....

SQL> connect / as sysdba

Connected to an idle instance

SQL> startup nomount

ASM Instance started

Total System Global Area 100663296

Fixed Size 777616

Variable Size 99885680

Database Buffers 0

Redo Buffers 0

ASM- Automatic File Management

Overview of configuration process (continued)

Creating the Disk Groups:

```
SQL>Create diskgroup asmgrp1
2   normal redundancy
3   failgroup channel1 disk
4   '/dev/raw/raw1
5   failgroup channel2 disk
6   '/dev/raw/raw2';
```

Diskgroup created.

The 3 levels of redundancy are:

External – Minimum requirement is 1 disk.

**Normal – Minimum requirement is 2 disks
– mirrored Available disk space 50% of
physical Disk.**

**High – Minimum requirement is 3 disks
triple mirrored. Available disk space is
33.33% of physical disk.**

Altering the Disk Groups:

```
SQL> alter diskgroup asmgrp1
2   add failgroup channel1 disk
3   '/dev/raw/raw3';
```

Diskgroup altered.

You can name the diskgroups as well

```
SQL> alter diskgroup asmgrp1
2   add failgroup channel2 disk
3   '/dev/raw/raw4'
4   name
5   disk73_02;
```

Diskgroup altered.

ASM- ASMLib (Linux)

Worth mentioning...

If installing ASM on Linux, you will need to install the Oracle ASMLib to increase I/O performance for particular Linux versions

You must download the following three packages, where version is the version of the ASM library driver, arch is the system architecture, and kernel is the version of the kernel that you are using:

- oracleasm-support-version.arch.rpm
- oracleasm-kernel-version.arch.rpm
- oracleasm-lib-version.arch.rpm

Install the drivers per the installation guide.....

You are then able to add the devices to the ASM instance and the system will be able to auto discover the devices.

This is only supported for certain Linux versions and doesn't work on some. It supports Redhat, but only RHAS 2.1 (it has a bug in 3.0). It also supports United Linux.

RMAN with ASM

RMAN and ASM –

- The use of aliases for file names will allow the DBA to find the correct file.
- In the case of incomplete recovery, the ASM can interact with RMAN, and an alias can be used to specify the appropriate datafile to delete, allowing RMAN to perform the necessary recovery procedure.

ASM – Views v\$ASM_DISKGROUP

```
Select *  
from v$asm_diskgroup
```

GROUP_NUMBER	NAME	SECTOR_SIZE	BLOCK_SIZE	ALLOCATION_UNIT_SIZE	STATE	TYPE	TOTAL_MB	FREE_MB
1	ASMGRP1	512	4096	1048576	MOUNTED	NORMAL	173166	137424

v\$asm_diskgroup – describes a disk group. The information available here is the:

Identifying group number – unique identification number

The group name – assigned group name

The sector size – 512 byte os level sectors

The block size – 4k block size

Allocation unit size

State – Current state of the disk group

Type – Redudancy level –(External, Normal, High)

Total MegaBytes – Its build size

Free MegaBytes – number of meg free

ASM – Views V\$ASM_CLIENT

```
select * from v$asm_client
```

GROUP_NUMBER	INSTANCE_N	DB_NAME	STATUS
----- 1	----- asmdb	----- asmdb	----- CONNECTED

v\$asm_client – identified databases using disk groups managed by the ASM instance

Unique group number
Instance Name Connected
Database Name Connected
Current connection status

ASM – Views V\$ASM_DISK

(continued)

V\$asm_disk

v\$asm_disk – all disks discovered by the ASM instance

GROUP_NUMBER	NUMBER		
DISK_NUMBER	NUMBER		
COMPOUND_INDEX	NUMBER		
INCARNATION	NUMBER		
MOUNT_STATUS	VARCHAR2(7)		
HEADER_STATUS	VARCHAR2(12)		
MODE_STATUS	VARCHAR2(7)		
STATE	VARCHAR2(8)		
REDUNDANCY	VARCHAR2(7)		
LIBRARY	VARCHAR2(64)		
TOTAL_MB	NUMBER		
FREE_MB	NUMBER	READS	NUMBER
NAME	VARCHAR2(30)	WRITES	NUMBER
FAILGROUP	VARCHAR2(30)	READ_ERRS	NUMBER
LABEL	VARCHAR2(31)	WRITE_ERRS	NUMBER
PATH	VARCHAR2(256)	READ_TIME	NUMBER
UDID	VARCHAR2(64)	WRITE_TIME	NUMBER
CREATE_DATE	DATE	BYTES_READ	NUMBER
MOUNT_DATE	DATE	BYTES_WRITTEN	NUMBER
REPAIR_TIMER	NUMBER		

ASM – Views V\$ASM_DISK

(continued)

v\$asm_disk – all disks discovered by the ASM instance

Select * from v\$asm_disk

DISK_NUMBER	MOUNT_STATUS	MOUNT_DATE	FAILGROUP
3	CACHED	24-APR-04	CHANNEL2
2	CACHED	24-APR-04	CHANNEL1
1	CACHED	24-APR-04	CHANNEL2
0	CACHED	24-APR-04	CHANNEL1

DISK_NUMBER	NAME	FAILGROUP	READ_ERRS	WRITE_ERRS
3	DISK73_02	CHANNEL2	0	0
2	ASMGRP1_0002	CHANNEL1	0	0
1	ASMGRP1_0001	CHANNEL2	0	0
0	ASMGRP1_0000	CHANNEL1	0	0

ASM – Views V\$ASM_FILE

v\$asm_file – every file in the disk group mounted by the ASM instance

```
SQL> desc v$asm_file
```

Name	Null?	Type
GROUP_NUMBER		NUMBER
FILE_NUMBER		NUMBER
COMPOUND_INDEX		NUMBER
INCARNATION		NUMBER
BLOCK_SIZE		NUMBER
BLOCKS		NUMBER
BYTES		NUMBER
SPACE		NUMBER
TYPE		VARCHAR2(64)
REDUNDANCY		VARCHAR2(6)
STRIPED		VARCHAR2(6)
CREATION_DATE		DATE
MODIFICATION_DATE		DATE

ASM – Views V\$ASM_FILE

v\$asm_file – every file in the disk group mounted by the ASM instance

Select * from v\$asm_file

GROUP_NUMBER	FILE_NUMBER	COMPOUND_INDEX	BYTES	STRIPE
-----	-----	-----	-----	-----
1	256	16777472	12173312	FINE
1	257	16777473	20972032	FINE
1	258	16777474	20972032	FINE
1	259	16777475	20972032	FINE
1	260	16777476	20972032	FINE
1	261	16777477	20972032	FINE
1	262	16777478	20972032	FINE
1	263	16777479	20972032	FINE
1	264	16777480	20972032	FINE
1	265	16777481	20972032	FINE
1	266	16777482	20972032	FINE

ASM – Views V\$ASM_TEMPLATE

v\$asm_template – one row for every template present in the disk group mounted by the instance

```
desc v$asm_template
```

Name	Null?	Type
GROUP_NUMBER		NUMBER
ENTRY_NUMBER		NUMBER
REDUNDANCY		VARCHAR2(6)
STRIPE		VARCHAR2(6)
SYSTEM		VARCHAR2(1)
NAME		VARCHAR2(30)

```
select * from v$asm_template
```

GROUP_NUMBER	ENTRY_NUMBER	REDUND	STRIPE	SYS	NAME
1	0	MIRROR	COARSE	Y	PARAMETERFILE
1	1	MIRROR	COARSE	Y	DUMPSET
1	2	MIRROR	FINE	Y	CONTROLFILE
1	3	MIRROR	COARSE	Y	ARCHIVELOG
1	4	MIRROR	FINE	Y	ONLINELOG
1	5	MIRROR	COARSE	Y	DATAFILE
1	6	MIRROR	COARSE	Y	TEMPFILE

ASM – Views V\$ASM_ALIAS

v\$asm_alias – contains one row for every alias present in every disk group mounted by the instance

desc v\$asm_alias

Name	Null?	Type
NAME		VARCHAR2(48)
GROUP_NUMBER		NUMBER
FILE_INCARNATION		NUMBER
ALIAS_INDEX		NUMBER
ALIAS_INCARNATION		NUMBER
PARENT_INDEX		NUMBER
REFERENCE_INDEX		NUMBER
ALIAS_DIRECTORY		VARCHAR2(1)
SYSTEM_CREATED		VARCHAR2(1)

ASM – Views V\$ASM_ALIAS

v\$asm_alias – contains one row for every alias present in every disk group mounted by the instance

```
select * from v$asm_alias
```

NAME	GROUP_NUMBER
-----	-----
TOOLS.366.1	1
system_01.dbf	1
undo_01.dbf	1
sysaux_01.dbf	1
TEMPFILE	1
TEMP.325.1	1
temp_01.dbf	1
arch	1
asmdb_1_38_524366675.arc	1
asmdb_1_37_524366675.arc	1
asmdb_1_39_524366675.arc	1

ASM – Views V\$ASM_OPERATION

v\$asm_operation – contains one row for every active ASM long running operation executing in the ASM instance

```
desc v$asm_operation
```

Name	Null?	Type
GROUP_NUMBER		NUMBER
OPERATION		VARCHAR2(5)
STATE		VARCHAR2(4)
POWER		NUMBER
ACTUAL		NUMBER
SOFAR		NUMBER
EST_WORK		NUMBER
EST_RATE		NUMBER
EST_MINUTES		NUMBER

D

Automatic File Management

We will discuss the options for 2 databases One with OMF and one without.....

Create a directory for each of the databases:

```
SQL> alter diskgroup asmgrp1 add directory '+asmgrp1/asmdb'  
diskgroup altered.
```

```
SQL> alter diskgroup asmgrp1 add directory '+asmgrp1/omfdb'  
diskgroup altered.
```



Automatic File Management

(continued)

```
----- begin -----  
*.audit_file_dest='/ora00/app/oracle/admin/asmdb/adump'  
*.audit_trail='db' *.background_dump_dest='/ora00/app/oracle/admin/asmdb/bdump'  
*.compatible='10.1.0'  
*.control_file_record_keep_time=10  
*.control_files='+asmgrp1/asmdb/controlasmdb.con','/ora01/oradata/asmdb/controlasmdb.con','/var/opt/oracle/oradata/asmdb/controlasmdb.con'  
*.core_dump_dest='/ora00/app/oracle/admin/asmdb/cdump'  
*.db_block_size=8192  
*.db_domain='tenure.com'  
*.db_file_multiblock_read_count=32  
*.db_files=200 *.db_name='asmdb'  
*.disk_asynch_io=true # ENABLED  
*.instance_name='asmdb'  
*.instance_type='rdbms'  
*.job_queue_processes=4  
*.log_archive_dest_1='location='+asmgrp1/asmdb/arch reopen=60'  
*.log_archive_format='asmdb_%t_%s_%r.arc'  
*.log_buffer=1048576# 1 MB *.log_checkpoint_interval=9999999 # CP AT LOG SWITCH ONLY  
*.log_checkpoint_timeout=0 # Infinite  
*.log_checkpoints_to_alert=true  
*.max_dump_file_size='2048'# 2048  
*.512 OS = 10MB  
*.open_cursors=100  
*.open_links=4
```

Automatic File Management

(continued)

Parameters continued

```
*.os_authent_prefix=""  
*.pga_aggregate_target=100M  
*.processes=300  
*.remote_login_passwordfile='exclusive'  
*.remote_os_authent=false  
*.resource_limit=true  
*.service_names='asmdb.achbill.com'  
*.sga_max_size=500M *.sga_target=400M  
*.sql_trace=false  
*.statistics_level='typical'  
*.timed_statistics=true  
*.undo_management='auto'  
*.undo_tablespace='undo'  
*.user_dump_dest='/ora00/app/oracle/admin/asmdb/udump'  
*.workarea_size_policy='auto'
```

Automatic File Management

(continued)

```
SQL> connect / as sysdba
Connected to an idle instance.
SQL> startup nomount
pfile=/ora00/app/oracle/admin/asmdb/pfile/initasmd
b.ora
```

ORACLE instance started.

Total System Global Area	524288000 bytes
Fixed Size	779660 bytes
Variable Size	28858484 bytes
Database Buffers	293601280 bytes
Redo Buffers	1048576 bytes

```
SQL> create database "asmdb"
2 maxinstances1
3 maxlogfiles 100
4 maxdatafiles 3000
5 maxlogmembers 5
6 character set "UTF8"
7 national character set "UTF8"
8 datafile '+asmgrp1/asmdb/system_01.dbf' size 325M
9 extent management local
10 sysaux
11 datafile '+asmgrp1/asmdb/sysaux_01.dbf' size 201M
12 default temporary tablespace temp
13 tempfile '+asmgrp1/asmdb/temp_01.dbf' size 1025M
14 uniform size 72K
15 undo tablespace undo
16 datafile '+asmgrp1/asmdb/undo_01.dbf' size 1025M
17 logfile
18 group 1 ('+asmgrp1/asmdb/redo0101.rdo',
19 '+asmgrp1/asmdb/redo0102.rdo') size 20M,
20 group 2 ('+asmgrp1/asmdb/redo0201.rdo',
21 '+asmgrp1/asmdb/redo0202.rdo') size 20M,
22 group 3 ('+asmgrp1/asmdb/redo0301.rdo',
23 '+asmgrp1/asmdb/redo0302.rdo') size 20M,
24 group 4 ('+asmgrp1/asmdb/redo0401.rdo',
25 '+asmgrp1/asmdb/redo0402.rdo') size 20M,
26 group 5 ('+asmgrp1/asmdb/redo0501.rdo',
27 '+asmgrp1/asmdb/redo0502.rdo') size 20M,
28 group 6 ('+asmgrp1/asmdb/redo0601.rdo',
29 '+asmgrp1/asmdb/redo0602.rdo') size 20M
30 /
Database created.
```

Automatic File Management

(continued)

SQL> list

1 select name, path, total_mb, free_mb

2* from v\$asm_disk

SQL> /

NAME	PATH	TOTAL_MB	FREE_MB
DISK73_02	/dev/raw/raw4	69397	56391
ASMGRP1_0002	/dev/raw/raw3	69397	56392
ASMGRP1_0001	/dev/raw/raw2	17186	13901
ASMGRP1_0000	/dev/raw/raw1	17186	13900

SQL> select * from v\$asm_client;

GROUP_NUMBER	INSTANCE_NAME	DB_NAME	STATUS
1	asmdb	asmdb	CONNECTED

You can see that the asmdb instance is connected to the ASM instance and using ASM resources.

Automatic File Management

(continued)

```
----- begin -----  
*.audit_file_dest='/ora00/app/oracle/admin/omfdb/adump'  
*.audit_trail='db'  
*.background_dump_dest='/ora00/app/oracle/admin/omfdb/bdump'  
*.compatible='10.1.0'  
*.control_file_record_keep_time=10  
*.control_files='+asmgrp1/omfdb/controlomfdb.con','/ora01/oradata/omfdb/controlomfdb.con','/var/opt/oracle/oradata/omfdb/controlomfdb.con'  
*.core_dump_dest='/ora00/app/oracle/admin/omfdb/cdump'  
*.db_block_size=8192  
*.db_create_file_dest='+asmgrp1'  
*.db_create_online_log_dest_1='+asmgrp1'  
*.db_domain='tenure.com'  
*.db_file_multiblock_read_count=32  
*.db_files=200  
*.db_name='omfdb'  
*.disk_asynch_io=true# ENABLED  
*.instance_name='omfdb'  
*.instance_type='rdbms'  
*.job_queue_processes=4  
*.log_archive_dest_1='location='+asmgrp1/omfdb/arch reopen=60'  
*.log_archive_format='omfdb_%t_%s_%r.arc'  
*.log_buffer=1048576# 1 MB  
*.log_checkpoint_interval=9999999# CP AT LOG SWITCH ONLY  
*.log_checkpoint_timeout=0# Infinite  
*.log_checkpoints_to_alert=true  
*.max_dump_file_size='2048'# 2048 * 512 OS = 10MB  
*.open_cursors=100  
*.open_links=4
```

Automatic File Management

(continued)

Parameter file continued

```
*.remote_login_passwordfile='exclusive'  
*.remote_os_authent=false  
*.resource_limit=true  
*.service_names='omfdb.achbill.com'  
*.sga_max_size=500M  
*.sga_target=400M  
*.sql_trace=false  
*.statistics_level='typical'  
*.timed_statistics=true  
*.user_dump_dest='/ora00/app/oracle/admin/omfdb/udump'  
*.workarea_size_policy='auto'
```

Automatic File Management

(continued)

```
SQL> connect / as sysdba
Connected to an idle instance.
SQL> startup nomount pfile=initomfdb.ora
ORACLE instance started.
```

```
Total System Global Area 524288000 bytes
Fixed Size                  779660 bytes
Variable Size               228858484 bytes
Database Buffers           293601280 bytes
Redo Buffers                1048576 bytes
SQL> create database omfdb;
```

Database created.

```
SQL>
```

ASM – Summary

ASM is indeed the future.

Good news for Disk management –

- By creating a meta-data instance to handle the interface between the physical disk and the database, we are simplifying the management of number of large disks that involve complexity of disk arrays.

Load balancing and I/O issues –

- The ASM instance takes care of load balancing for you as you add or remove disks from the disk array.

Thank you

You can reach us:

Dallas Headquarters:

- Tenure
222 West Las Colinas Blvd. Suite 1047
East Tower
Irving, TX 75039
- Phone: +1 (972)-432-1000

Special thanks to :

Byron Pearce bepearce@tenure.com

Tom Wood twood@tenure.com

Questions.....

