

ORACLE®



ORACLE®

**Top 10 Things You Always Wanted to Know About Automatic
Storage Management
But Were Afraid to Ask**

Nitin Vengurlekar
RAC/ASM Engineering
RacPack

Agenda

- ASM Overview – 2 minute tour
- Top Ten Questions
 - ASM Architecture
 - ASM Configuration & Performance
 - Backups
 - Migration
 - 3rd Party Software
- Summary



ORACLE

ASM Architecture - Review

- Automatic Storage Management (ASM) instance
 - Instance that manages the diskgroup metadata
- Disk Groups
 - Logical grouping of disks
 - Determines file mirroring options
- ASM Disks
 - LUNs presented to ASM
- ASM Files
 - Files that are stored in ASM disk groups are called ASM files, this includes database files

ORACLE

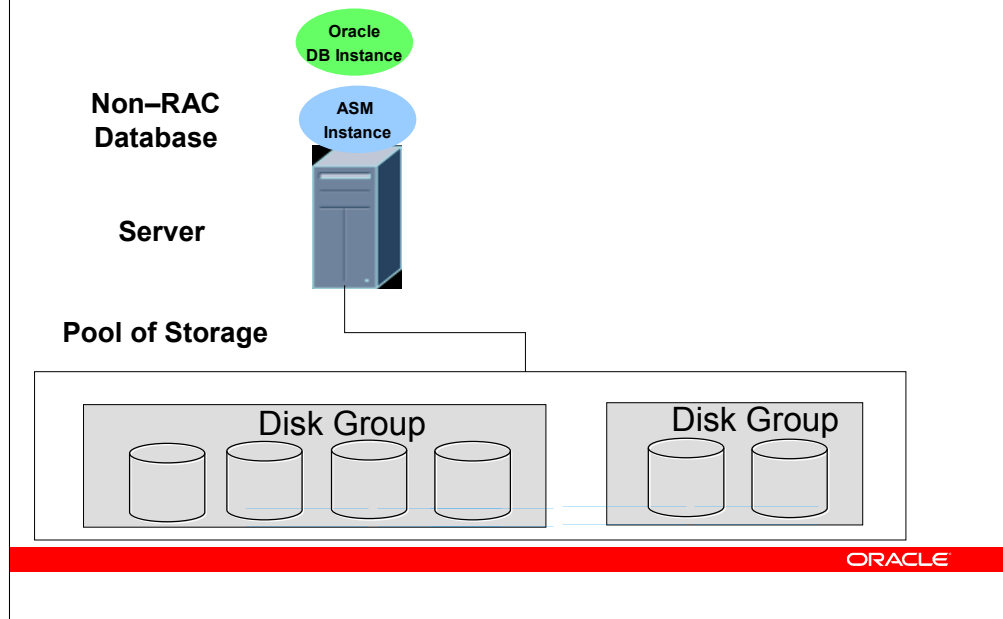
In Oracle Database 10g/11g there are two types of instances: database and ASM instances. The ASM instance, which is generally named +ASM, is started with the INSTANCE_TYPE=ASM init.ora parameter. This parameter, when set, signals the Oracle initialization routine to start an ASM instance and not a standard database instance. Unlike the standard database instance, the ASM instance contains no physical files; such as logfiles, controlfiles or datafiles, and only requires a few init.ora parameters for startup.

Upon startup, an ASM instance will spawn all the basic background processes, plus some new ones that are specific to the operation of ASM. The STARTUP clauses for ASM instances are similar to those for database instances. For example, RESTRICT prevents database instances from connecting to this ASM instance. NOMOUNT starts up an ASM instance without mounting any disk group. MOUNT option simply mounts all defined diskgroups
For RAC configurations, the ASM SID is +ASMx instance, where x represents the instance number.

A disk group consists of multiple disks and is the fundamental object that ASM manages. Each disk group contains the metadata that is required for the management of space in the disk group. The ASM instance manages the metadata about the files in a Disk Group in the same way that a file system manages metadata about its files. However, the vast majority of I/O operations do not pass through the ASM instance. In a moment we will look at how file I/O works with respect to the ASM instance.

ASM Architecture

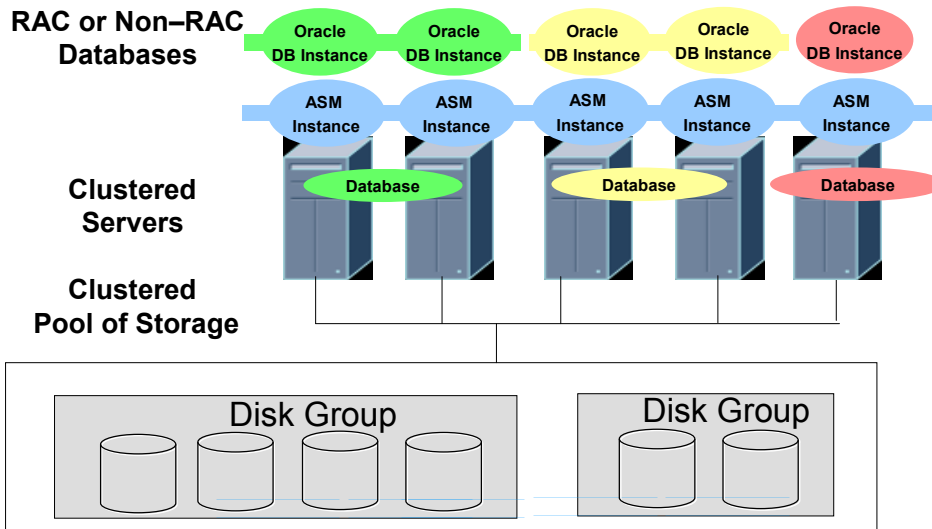
A Brief Overview



As all DBAs knows, the database software instantiation on a server is called a database instance. ASM's nomenclature is the same in that there is an ASM instance. Like the database, ASM has a shared memory footprint called the SGA. The database and ASM instances communicate with each other through high-performance O/S primitives. There is one ASM instance per server. Even when there are several database instances on a single server, each database shares a single ASM instance on that server.

ASM Process Architecture

A Brief Overview



ORACLE

- In a RAC configuration, each server in the cluster has an ASM instance. Each of these instances communicates amongst themselves in the management and presentation of the file system.
- Each database instance coordinates file access through the ASM instance on the server for which it operates.
- Lastly, each ASM instance gathers performance related data from its associated database instance on its server.

Top 10 ASM Questions
ASM Architecture



ORACLE

Top 10 ASM Questions

Q. What init.ora parameters does a user need to configure for ASM instances?

A. The default parameter settings work perfectly for ASM. The only parameters needed for 11g ASM:

- PROCESSES*
- ASM_DISKSTRING*
- ASM_DISKGROUPS
- INSTANCE_TYPE

ORACLE

•ASM is a very passive instance in that it doesn't have a lot concurrent transactions or queries. So the memory footprint is quite small.

•Even if you have 20 dbs connected to ASM , the ASM SGA does not need to change. This is b/c the ASM metadata is not directly tied to the number of clients

•The 11g MEMORY_TARGET (DEFAULT VALUE) will be more than sufficient.

•The processes parameter may need to be modified. Use the formula to determine the approp value:

```
processes = 40 + (10 + [max number of concurrent database file  
creations, and file extend operations possible])*n
```

Where n is the number of databases connecting to ASM (ASM clients).

The source of concurrent file creations can be any of the following:

- Several concurrent create tablespace commands
- Creation of a Partitioned table with several tablespaces creations
- RMAN backup channels
- Concurrent archive logfile creations

Top 10 ASM Questions

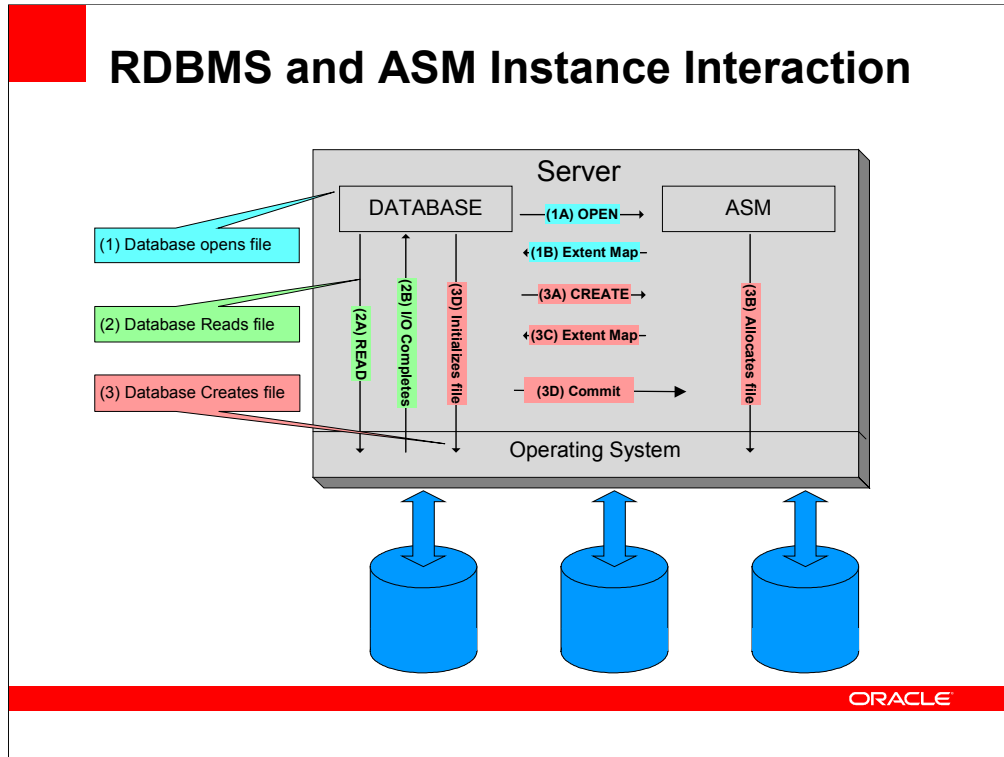
Q. How does the database interact with the ASM instance and how do I make ASM go faster?

A. ASM is not in the I/O path so ASM does not impede the database file access. Since the RDBMS instance is performing raw I/O, the I/O is as fast as possible.

ORACLE

- Cover ASM instance architecture
- Cover ASM-Communication via ASMB
- The database communicates with ASM instance using the ASMB (umblicus process) process. Once the database obtains the necessary extents from extent map, all database IO going forward is processed through by the database processes, bypassing ASM. Thus we say ASM is not really in the IO path. So, the question how do we make ASM go faster.....you don't have to.

RDBMS and ASM Instance Interaction



1A. Database issues open of a database file

1B. ASM sends the extent map for the file to database instance. Starting with 11g, the RDBMS only receives first 60 extents the remaining extents in the extent map are paged in on demand, providing a faster open

2A/2B. Database now reads directly from disk

3A.RDBMS foreground initiates a create tablespace for example

3B. ASM does the allocation for its essentially reserving the allocation units for the file creation

3C. Once allocation phase is done, the extent map is sent to the RDBMS

3D. The RDBMS initialization phase kicks in. In this phase the initializes all the reserved AUs

3E. If file creation is successful, then the RDBMS commits the file creation

Going forward all I/Os are done by the RDBMS.

Top 10 ASM Questions

Q. Do I need to define the RDBMS FILESYSTEMIO_OPTIONS parameter when I use ASM?

A. No. The RDBMS does I/O directly to the raw disk devices, the FILESYSTEMIO_OPTIONS parameter is only for filesystems.

ORACLE

A. Review what the use of FILESYSTEMIO_OPTIONS parameter; essentially FILESYSTEMIO_OPTIONS is used for filesystem/block storage.

This parameter controls which IO options are used. The value may be any of the following:

*asynch - This allows asynchronous IO to be used where supported by the OS.

*directIO - This allows directIO to be used where supported by the OS.

Direct IO bypasses any Unix buffer cache. *setall - Enables both ASYNC and DIRECT IO. "none" - This disables ASYNC IO and DIRECT IO so that Oracle uses normal synchronous writes, without any direct io options.

A. RDBMS does raw IO against the ASM disks, so need for FILESYSTEMIO_OPTIONS parameter. The only parameter that needs to be set is [disk_asyncio=true](#), which is true by default. If using ASMLIB then even the [disk_async](#) does not need to be set.

ASM is also supported for NFS files as ASM disks. In such cases, the required NFS mount options eliminate the need to set FILESYSTEMIO_OPTIONS.

Top 10 ASM Questions
ASM Configuration



ORACLE

Top 10 ASM Questions

Q. I read in the ASM Best Practices paper that Oracle recommends two diskgroups. Why?

A. Oracle recommends two diskgroups to provide a balance of manageability, utilization, and performance.

ORACLE

To reduce the complexity of managing ASM and its diskgroups, Oracle recommends that generally no more than two diskgroups be maintained and managed per RAC cluster or single ASM instance

oDatabase work area: This is where active database files such as datafiles, control files, online redo logs, and change tracking files used in incremental backups are stored. This location is indicated by DB_CREATE_FILE_DEST.

oFlash recovery area: Where recovery-related files are created, such as multiplexed copies of the current control file and online redo logs, archived redo logs, backup sets, and flashback log files. This location is indicated by DB-RECOVERY_FILE_DEST.

•Having one DATA container means only place to store all your database files, and obviates the need to juggle around datafiles or having to decide where to place a new tablespace.

By having one container for all your files also means better storage utilization. Making the IT director very happy. If more storage capacity or IO capacity is needed, just add an ASM disk....all online activities.

You have to ensure that this storage pool container houses enough spindles to accommodate the IO rate of all the database objects

Bottom line, one container == one pool manage, monitor, and track

Note however, that additional diskgroups may be added to support tiered storage classes in Information Lifecycle Management (ILM) or Hierarchical Storage Management (HSM) deployments

Top 10 ASM Questions

Q. We have a 16 TB database. I'm curious about the number of disk groups we should use; e.g. 1 large disk group, a couple of disk groups, or otherwise?

A. For VLDBs you will probably end up with different storage tiers; e.g with some of our large customers they have Tier1 (RAID10 FC), Tier2 (RAID5 FC), Tier3 (SATA), etc. Each one of these is mapped to a diskgroup.

ORACLE

These custs mapped certain tablespaces to specific tiers; eg, system/rollback/syaux and latency sensitive tablespaces in Tier1, and not as IO critical on Tier2, etc.

For 10g VLDBs its best to set an AU size of 16MB, this is more for metadata space efficiency than for performance. The 16MB recommendation is only necessary if the diskgroup is going to be used by 10g databases. In 11g we introduced variable size extents to solve the metadata problem. This requires compatible.rdbms & compatible.asm to be set to 11.1.0.0. With 11g you should set your AU size to the largest I/O that you wish to issue for sequential access (other parameters need to be set to increase the I/O size issued by Oracle). For random small I/Os the AU size does not matter very much as long as every file is broken into many more extents than there are disks.

Top 10 ASM Questions

Q. We have a new app and don't know our access pattern, but assuming mostly sequential access, what size would be a good AU fit?

A. For 11g ASM/RDBMS it is recommended to use 4MB ASM AU for disk groups.
See Metalink Note [810484.1](#)

ORACLE

For all 11g ASM/DB users, it best to create a disk group using 4 MB ASM AU size. Metalink Note 810484.1 covers this

Top 10 ASM Questions

Q. Would it be better to use BIGFILE tablespaces, or standard tablespaces for ASM?

A. The use of Bigfile tablespaces has no bearing on ASM (or vice versa). In fact most database object related decisions are transparent to ASM.

ORACLE

Nevertheless, Bigfile tablespaces benefits:

Fewer datafiles - which means faster database open (fewer files to open),

Faster checkpoints, as well fewer files to manage. But you'll need careful consideration for backup/recovery of these large datafiles.

Top 10 ASM Questions

Q. What is the best LUN size for ASM

A. There is no best size! In most cases the storage team will dictate to you based on their standardized LUN size. The ASM administrator merely has to communicate the ASM Best Practices and application characteristics to storage folks :

- Need equally sized / performance LUNs
- Minimum of 4 LUNs
- The capacity requirement
- The workload characteristic (random r/w, sequential r/w) & any response time SLA

Using this info , and their standards, the storage folks should build a nice LUN group set for you.

ORACLE

In most cases the storage team will dictate to you what the standardized LUN size is. This is based on several factors, including RAID LUN set builds (concatenated, striped, hypers, etc..). Having too many LUNs elongates boot time and is it very hard to manage On the flip side, having too few LUNs makes array cache management difficult to control and creates un-manageable large LUNs (which are difficult to expand).

The ASM administrator merely has to communicate to SA/storage folks that you need equally sized/performance LUNs and what the capacity requirement is, say 10TB. Using this info, the workload characteristic (random r/w, sequential r/w), and their standards, the storage folks should build a nice LUN group set for you

Having too many LUNs elongates boot time and is it very hard to manage (zoning, provisioning, masking, etc..)....there's a \$/LUN barometer!

Top 10 ASM Questions

Q. In 11g RAC we want to separate ASM admins from DBAs and create different users and groups. How do we set this up?

A. For clarification

- Separate Oracle Home for ASM and RDBMS.
- RDBMS instance connects to ASM using OSDBA group of the ASM instance. Thus, software owner for each RDBMS instance connecting to ASM must be a member of ASM's OSDBA group.
- Choose a different OSDBA group for ASM instance (asmdba) than for RDBMS instance (dba)
- In 11g, ASM administrator has to be member of a separate SYSASM group to separate ASM Admin and DBAs.

ORACLE

Operating system authentication using membership in the group or groups designated

as OSDBA, OSOPER, and OSASM is valid on all Oracle platforms.

A typical deployment could be as follows:

ASM administrator:

User : asm

Group: oinstall, asmdba(OSDBA), asmadmin(OSASM)

Database administrator:

User : oracle

Group: oinstall, asmdba(OSDBA of ASM), dba(OSDBA)

ASM disk ownership : asm:oinstall

Remember that Database instance connects to ASM instance as sysdba. The user id the database instance runs as needs to be the OSDBA group of the ASM instance.

Top 10 ASM Questions

A typical deployment could be as follows:

ASM administrator:

User : asm

Group: oinstall, asmdba(OSDBA), asmadmin(OSASM)

Database administrator:

User : oracle

Group: oinstall, asmdba(OSDBA of ASM), dba(OSDBA)

ORACLE

A typical deployment could be as follows:

ASM administrator:

User : asm

Group: oinstall, asmdba(OSDBA), asmadmin(OSASM)

Database administrator:

User : oracle

Group: oinstall, asmdba(OSDBA of ASM), dba(OSDBA)

Top 10 ASM Questions

Q. Can my RDBMS and ASM instances run different versions?

A. Yes. ASM can be at a higher version or at lower version than its client databases. There's two components of compatibility:

- Software compatibility
- Diskgroup compatibility attributes:
 - compatible.asm
 - compatible.rdbms

ORACLE

This is a diskgroup level change and not an instance level change...no rolling upgrade here!

Disk Group Compatibility Example

- Start with 10g ASM and RDBMS
- Upgrade ASM to 11g
- Advance compatible.asm
 - ALTER DISKGROUP data
SET ATTRIBUTE
'compatible.asm' = '11.1.0.7.0'
- 10g RDBMS instances are still supported
- 10g ASM instance can no longer mount the disk group

ORACLE

Disk Group Compatibility Example

- Upgrade RDBMS to 11g
- In the RDBMS instance set initialization parameter
 - `compatible = 11.0`
- Advance `compatible.rdbms`
 - `ALTER DISKGROUP data`
`SET ATTRIBUTE`
`'compatible.rdbms' = '11.1.0.7.0'`
- New capabilities enabled
 - Variable size extents
 - Fast mirror resync
 - Preferred read
 - AUs > 16MB
- 10g RDBMS instances can no longer access the disk group

ORACLE

Disk Group Compatibility Example

- Compatibility may be set during disk group creation
 - ```
CREATE DISKGROUP data
 DISK '/dev/sdd[bcd]1'
 ATTRIBUTE
 'compatible.asm' = '11.1.0.7.0',
 'compatible.rdbms' = '11.1.0.7.0',
 'au_size' = '4M'
```
- `compatible.asm` **and** `compatible.rdbms` **cannot be reversed**

ORACLE

## Top 10 ASM Questions

Q. Where do I run my database listener from; i.e., ASM HOME or DB HOME?

A. It is recommended to run the listener from the ASM HOME. This is particularly important for RAC env, since the listener is a node-level resource. In this config, you can create additional [user] listeners from the database homes as needed.

ORACLE

- Allows registering multiple databases on the node to register with the listener without being tied to a specific database home
- From configuration tool standpoint (netca), we promote best practice of creating one listener per node with node name suffix (that is registered with CRS) and subsequent tools that create/upgrade databases will register instances to that listener. One can always create multiple listeners in different homes and use'em but that would complicate the configuration

# Top 10 ASM Questions

## Backups



ORACLE

## Top 10 ASM Questions

Q. How do I backup my ASM instance?

A. Not applicable! ASM has no files to backup

ORACLE

**Unlike the database, ASM does require a controlfile type structure or any other external metadata to bootstrap itself. All the data ASM needs to startup is on on-disk structures (disk headers and other disk group metadata).**

**A Disk Group is the fundamental object managed by ASM. It is composed of multiple ASM disks. Each Disk Group is self-describing, like a standard file system. All the metadata about the usage of the space in the disk group is completely contained within the disk group. If ASM can find all the disks in a disk group it can provide access to the disk group without any additional metadata**

## Top 10 ASM Questions

Q. When should I use RMAN and when should I use ASMCMD copy?

A. RMAN is the recommended and most complete and flexible method to backup and transport database files in ASM.

ASMCMD copy is good for copying single files

- Supports all Oracle file types
- Can be used to instantiate a Data Guard environment
- Does not update the controlfile
- Does not create OMF files

ORACLE

**RMAN is the most complete and versatile means to backup databases stored in ASM.**

**However, many customers use BCV/split-mirrors as backups for ASM based databases. Many combine BCV mirrors with RMAN backup of the mirrors. Why would you want to do that? RMAN ensures the integrity of the database data blocks by running sanity checks as it backs up the blocks**

**Now most of you are wondering about the 11g asmcmd copy command, and how that fits in here. asmcmd cp is not intended to do wholesale backups (plus you'll have to put the database in hot backup).**

**In 10g the possible ways to migrate - DBMS\_FILE\_TRANSFER, rman (copy vs. backup), or XMLDB FTP**

**In 11g, we introduced the asmcmd copy command. Key point here is that copy files out is great for:**

1. archive logs
2. Controlfiles
3. Datafiles for debugging
4. Dumpsets (can be done across platforms)

**Copying files in:**

**TTS**

## ASMCMD Copy

```
ASMCMD> ls
+fra/dumpsets/expdp_5_5.dat

ASMCMD> cp expdp_5_5.dat sys@rac1.orcl1:+DATA/dumpsets/exp
pdp_5_5.dat
source +fra/dumpsets/expdp_5_5.dat
target +DATA/dumpsets/expdp_5_5.dat
copying file(s)...
file, +DATA/dumpsets/expdp_5_5.dat,
copy committed.
```

ORACLE

# Top 10 ASM Questions

## Migration



ORACLE

## Top 10 ASM Questions

Q. I'm going to do add disks to my ASM diskgroup, how long will this rebalance take?

A. Rebalance time is heavily driven by the three items:

- Amount of data currently in the diskgroup
- IO bandwidth available on the server
- ASM\_POWER\_LIMIT or Rebalance Power Level

ORACLE

Use `v$asm_operation`

## Top 10 ASM Questions

Q. We are migrating to a new storage array. How do I move my ASM database from storage A to storage B?

A. Given that the new and old storage are both visible to ASM, simply add the new disks to the ASM disk group and drop the old disks. ASM rebalance will migrate data online. Note 428681.1 covers how to move OCR/Voting disks to the new storage array

ORACLE


**If this is a RAC environment, the Note 428681.1 covers how to move OCR/Voting disks to the new storage array**

## Top 10 ASM Questions

```
ASM_SQL> alter diskgroup DATA
drop disk
 data_legacy1, data_legacy2,
 data_legacy3
add disk
 '/dev/sddb1', '/dev/sddc1',
 '/dev/sddd1';
```

ORACLE

## ASM Rebalancing




Disk Group DATA (legacy disks)

- Automatic online rebalance whenever storage configuration changes
- Only move data proportional to storage added
- No need for manual I/O tuning

ORACLE

## ASM Rebalancing

- Automatic online rebalance whenever storage configuration changes
- Online migration to new storage




Disk Group DATA

ORACLE

## ASM Rebalancing

- Automatic online rebalance whenever storage configuration changes
- Online migration to new storage




Disk Group DATA

ORACLE

## ASM Rebalancing

- Automatic online rebalance whenever storage configuration changes
- Online migration to new storage




Disk Group DATA

ORACLE

## ASM Rebalancing

- Automatic online rebalance whenever storage configuration changes
- Online migration to new storage



Disk Group DATA (new disks)

ORACLE

## Top 10 ASM Questions

Q. Is it possible to unplug an ASM disk group from one platform and plug into a server on another platform (for example, from Solaris to Linux)?

A. No. Cross-platform disk group migration not supported. To move datafiles between endian-ness platforms, you need to use XTTS, Datapump or Streams.

ORACLE

The first problem that you run into here is that Solaris and Linux format their disks differently. Solaris and Linux do not recognize each other's partitions, etc.

ASM does track the endian-ness of its data. However, currently, the ASM code does not handle disk groups whose endian-ness does not match that of the ASM binary.

Experiments have been done to show that ASM disk groups can be migrated from platforms that share a common format and endian-ness (i.e. Windows to Linux), but this functionality is not currently officially supported because is not regularly tested yet.

The following links show how to migrate across platforms

[http://download-west.oracle.com/docs/cd/B19306\\_01/server.102/b25159/outage.htm#CA CFFIDD](http://download-west.oracle.com/docs/cd/B19306_01/server.102/b25159/outage.htm#CA CFFIDD)

[http://www.oracle.com/technology/deploy/availability/pdf/MAA\\_WP\\_10gR2\\_PlatformMigrationTTS.pdf](http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_PlatformMigrationTTS.pdf)

## **Top 10 ASM Questions**

### **3<sup>rd</sup> Party Software**



ORACLE

## Top 10 ASM Questions

Q. How does ASM work with multipathing software?

A. It works great! Multipathing software is at a layer lower than ASM, and thus is transparent.

You may need to adjust `ASM_DISKSTRING` to specify only the path to the multipathing pseudo devices.

ORACLE

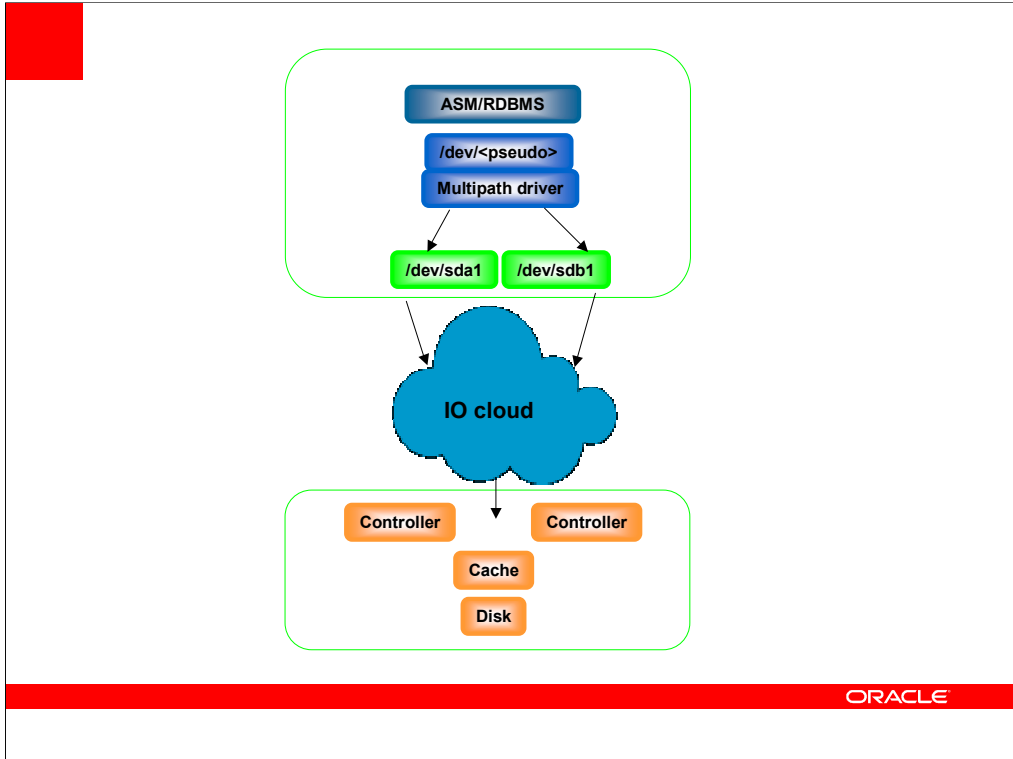
Multipathing tools provides the following benefits:

- o Provide a single block device interface for a multi-pathed LUN
- o Detect any component failures in the I/O path; e.g., fabric port, channel adapter, or HBA.
- o When a loss of path occurs, ensure that I/Os are re-routed to the available paths, with no process disruption.
- o Reconfigure the multipaths automatically when events occur.
- o Ensure that failed paths get revalidated as soon as possible and provide auto-failback capabilities.
- o Configure the multi-paths to maximize performance using various load balancing methods; e.g., round robin, least I/Os queued, or least service time.

When a given disk has several paths defined, each one will be presented as a unique path name at the OS level; e.g.; `/dev/rdisk/c3t19d1s4` and `/dev/rdisk/c7t22d1s4` could be pointing to same disk device. ASM, however, can only tolerate the discovery of one unique device path per disk. For example, if the `asm_diskstring` is `'/dev/rdisk/*'`, then several paths to the same device will be discovered, and ASM will produce an error message stating this. When using a multipath driver, which sits above this SCSI-block layer, the driver will generally produce a pseudo device that virtualizes the sub-paths. For example, in the case of EMC's PowerPath, you can use the following `asm_diskstring` setting of `'/dev/rdisk/emcpower*'`. When I/O is issued to this disk device, the multipath driver will intercept it and provide the necessary load balancing to the underlying subpaths.

As long as ASM can open/read/write to the multipathing pseudo device, it should work. Most all MP products are known to work w/ ASM. But remember ASM does not certify MP products, though we have a list products that work w/ ASM, this is more of a guide of what's available by platform/OS.

Examples of multi-pathing software include EMC PowerPath, Veritas DMP, Sun Traffic



## Top 10 ASM Questions

Q. Is ASM constantly rebalancing to manage “hot spots”?

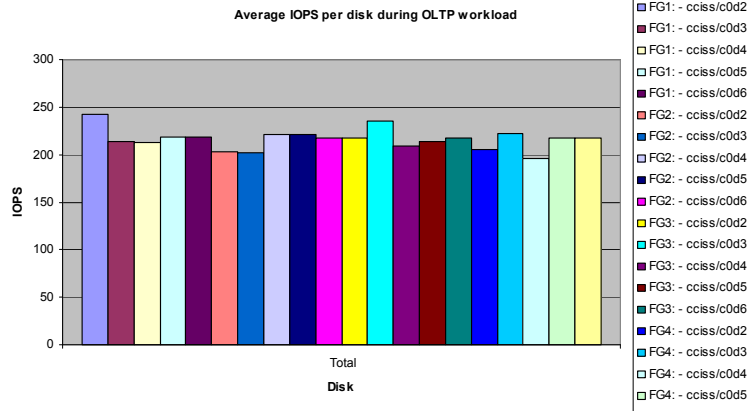
A. No...No...Nope!!

ORACLE

**Bad rumor. ASM provides even distribution of extents across all disks in a disk group. Since each disk will equal number of extents, no single disk will be hotter than another. Thus the answer *NO*, ASM does not dynamically move hot spots, because hot spots simply do not occur in ASM configurations.**

**Rebalance only occurs on storage configuration changes (e.g. add, drop, or resize disks).**

## I/O Distribution



- ASM spreads file extents evenly across all disks in disk group
  - Since ASM distributes extents evenly, there are no hot spots

ORACLE

As indicated, ASM implements the policy of S.A.M.E. that stripes and mirrors files across all the disks in a Disk Group. If the disks are highly reliable as the case may be with a high-end array, mirroring can be optionally disabled for a particular Disk Group. This policy of striping and mirroring across all disks in a disk group greatly simplifies storage management and provides a configuration of balanced performance.

## Key Value Propositions

- **Manageability**
  - Simple provisioning
  - Storage Array migration
  - VM/FS co-existence
  - SQL, EM, Command line
  - Consolidation
  - Self-tuning
- **Performance**
  - Distribute load across all available storage
  - No ASM code in data path
- **Availability**
  - Automatic mirror rebuild
  - Automatic bad block correction
  - Rolling upgrades
  - Online patches
  - RAC and clusterware support
- **Cost Savings**
  - Shared storage pool
  - Just-in-Time provisioning
  - No license fees
  - No support fees

ORACLE

## Summary:

- ASM requires very few parameters to run
- ASM based databases inherently leverage raw disk performance
- No additional database parameters needed to support ASM
- Mixed ASM-database version support
- Facilitates online storage changes
- RMAN recommended for backing up ASM based databases
- Spreads I/O evenly across all disks to maximize performance and eliminates hot spots

ORACLE

ASM provides filesystem and volume manager capabilities built into the Oracle database kernel. With this capability, ASM simplifies storage management tasks, such as creating/laying out databases and disk space management. Since ASM allows disk management to be done using familiar create/alter/drop SQL statements, DBAs do not need to learn a new skill set or make crucial decisions on provisioning.

The following are some key benefits of ASM:

- oASM spreads I/O evenly across all available disk drives to prevent hot spots and maximize performance.
- oASM eliminates the need for over provisioning and maximizes storage resource utilization facilitating database consolidation.
- oInherent large file support.
- oPerforms automatic online redistribution after the incremental addition or removal of storage capacity.
- oMaintains redundant copies of data to provide high availability, or leverages 3<sup>rd</sup> party RAID functionality.
- oSupports Oracle Database as well as Oracle Real Application Clusters (RAC).
- oCapable of leveraging 3<sup>rd</sup> party multipathing technologies.
- oFor simplicity and easier migration to ASM, an Oracle database can contain ASM and non-ASM files. Any new files can be created as ASM files whilst existing files can also be migrated to ASM.
- oRMAN commands enable non-ASM managed files to be relocated to an ASM disk group.
- oEnterprise Manager Database Control or Grid Control can be used to manage ASM disk and file activities.

## ASM Collateral and Content

<http://www.oracle.com/technology/asm>

- ASM 11g New Features
- ASM Best Practices
- ASM vendor papers
- ASM-RAC Customer Case Studies

ORACLE

**Top 10 ASM Questions**  
**Extra credit questions**



ORACLE

## Top 10 ASM Questions

Q. Is ASMLIB required on Linux systems and are there any benefits to using it?

A. ASMLIB is not required to run ASM, but it is certainly recommended.

ASMLIB has following benefits:

- Simplified disk discovery
- Persistent disk names
- Efficient use of system resources

ORACLE

- o Reduced Overhead

**ASMLIB provides the capability for a process (RBAL) to perform a global open/close on the disks that are being dropped or closed.**

**This reduces the number of open file descriptors on the system, thus reduces the probability of running out of global file descriptors on the system. Also, the open and close operations are reduced, ensuring orderly cleanup of file descriptors when storage configurations changes occur.**

**A side benefit of the aforementioned items is a faster startup of the database.**

- o Disk Management and Discovery

**With ASMLib the ASM disk name is automatically taken from the name given it by the administrative tool. This simplifies adding disks and correlating OS names with ASM names, as well as eliminates erroneous disk management activities since disks are already pre-named.**

**The default discovery string for ASM is NULL, however, if ASMLIB is used, the ASMLIB default string replaces the NULL string, making disk discovery much more straightforward. Note, disk discovery has been one of the big challenges for administrators.**

**The ASMLib permissions are persistent across reboot and in the event of major/minor number changes**

**In RAC environments, disk identification and discovery as simply as single instance. Once the disks are labeled on one node, the other clustered nodes simply use the default disk discovery string, and discovery is seamless.**

- o No requirement to setup raw links

## Top 10 ASM Questions

Q. Is it possible to do rolling upgrades on ASMLIB in a RAC configuration

A. ASMLIB is independent of Oracle Clusterware and Oracle Database, and thus can be upgraded on its own

ORACLE

Upgrading ASMLIB on a given node will require that ASMLIB be disabled/stop, will require the database and ASM to also be shutdown on that node. Once the ASMLIB is upgraded then the stack can be restarted.

**ORACLE®**