

Dell EMC SC Series Storage with Oracle ASM

Abstract

This document provides best practices and configuration options for Dell EMC™ SC Series storage in Fibre Channel environments when deploying Oracle®.

April 2018

Revisions

Date	Description
May 2011	Initial release
April 2012	Content and format change
January 2016	Content and format change
April 2018	Content and format change

Acknowledgements

Author: Mark Tomczik

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2011–2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	4
1 Introduction.....	5
1.1 How to use this document	5
1.2 Objective.....	5
1.3 Audience.....	5
1.4 Prerequisites.....	6
2 Technology overview.....	7
2.1 SC Series storage	7
2.2 Oracle technology.....	7
2.3 Oracle ASM features	7
2.3.1 Oracle Flex ASM features	8
2.3.2 Oracle ASMLib features	10
2.3.3 Benefits of ASM, Flex ASM, and ASMLib	10
3 Components of ASM	11
3.1 ASM instance.....	11
3.2 ASMLib	11
3.3 Disk groups.....	11
3.3.1 ASM disk group redundancy	12
4 ASM best practices with SC Series storage.....	13
5 Conclusion.....	15
A Flex ASM client query	16
B Other useful commands for an ASM environment	18
C 512 byte and 4K byte sectors.....	19
D Technical support and resources	20
D.1 Related resources	20

Executive summary

One key IT and business challenge is to manage and provision storage to accommodate the increasingly demanding requirements of a growing Oracle® database. Protecting a database and ensuring its integrity and availability is also mandatory for business success.

When designing the physical structure of a database, administrators must consider different storage configuration options. Storage solutions must facilitate high performance I/O and provide protection against failure of storage hardware such as disks and host bus adapters (HBAs). Growing and changing workloads require a dynamic storage configuration that provides a way to automate storage-related tasks, reducing the risk of human error.

The Oracle Automatic Storage Management (ASM) feature simplifies storage provisioning and data layout challenges. In combination with ASM, the Dell EMC™ SC Series thin provisioning and dynamic block architecture software allow Oracle data to be easily accessible, provide a data lifecycle management structure, and enable dynamic capacity in less time than other traditional storage solutions. The intent of this document is to help readers understand best practices of SC Series storage in Fibre Channel environments when deploying Oracle ASM or Flex ASM. This helps to deliver a cost-effective alternative that meets the performance and fault-tolerant storage requirements of Oracle deployments.

1 Introduction

Harnessing the power of Oracle ASM with the intelligence of SC Series storage yields a reliable, stable, and dynamic infrastructure that can be protected with very little effort. The following sections cover general best practices for using Oracle ASM with SC Series storage. For additional information on using Oracle and SC Series storage, some of which is applicable to ASM, it is strongly recommended that readers review the best practices document, [Dell EMC SC Series Arrays and Oracle](#).

1.1 How to use this document

The best practices in this document are generally applicable to other Linux® and software versions, but the commands may have to be modified to fit other versions. Also, the best practices and recommendations presented in this document should be evaluated and adjusted accordingly for each environment.

This document is not intended to be an exact step-by-step configuration guide. It should not be considered an exhaustive or authoritative source on any component discussed. Performance tuning of components is not in scope of this document. Actual configuration will vary based on individual business, application, infrastructure, and compatibility requirements, as well as Linux and software versions.

For information on SC Series arrays, refer to the *Dell Storage Center System Manager Administrator's Guide* available at the [SC Series customer portal](#) (login required). For detailed information on Oracle, see the [My Oracle Support](#) or [Oracle Help Center](#) sites. Additional resources are listed in appendix D.

1.2 Objective

Although this document is not a comprehensive user guide, it provides general best practices to combine Oracle ASM and Oracle Flex ASM with SC Series storage. After reviewing this document, reader will understand possible storage solutions and best practices to address Oracle database storage needs.

Note that the suggestions and solutions presented in this document should be evaluated and adjusted for individual environments. Dell EMC advises customers to understand the type of storage needs required by the business before designing and implementing a solution.

For detailed information on any part of the Oracle installation and configuration used in this paper, refer to the information provided at the My Oracle Support site. Oracle ASM tuning, converting between standard legacy ASM deployments and Flex ASM deployments, Flex clusters, and Oracle proxy ASM Instance are beyond the scope of this document. See the references listed in appendix D for additional information.

1.3 Audience

This document is intended for database administrators (DBAs), system administrators, storage administrators, and architects seeking to design, implement, and maintain robust and efficient database and storage systems for Oracle database solutions.

1.4 Prerequisites

Readers should have formal training or advanced working knowledge of the following:

- [Dell EMC SC Series with Red Hat Enterprise Linux \(RHEL\) 6x](#) best practices
- [Dell EMC SC Series Arrays and Oracle](#) best practices
- RAID, Fibre Channel, Serial-Attached SCSI (SAS), LUNs, and multipathing
- General understanding of SAN technologies
- Dell EMC server architecture, configuration, and administration
- Operating and configuring SC Series storage arrays
- Operating Dell Storage Manager (DSM)
- IPv4 network concepts and administration
- Familiarity of and installing, configuring, and administration of Red Hat® and Oracle Linux
- Installing, configuring, and administering Oracle 11g and 12c databases, Oracle ASM, and Oracle 11g and 12c Grid Infrastructure
- Oracle Real Application Cluster (RAC) and single-instance database

2 Technology overview

2.1 SC Series storage

SC Series storage provides features that can be used in an Oracle database environment to create a robust storage solution and assist DBAs with data tiering, database performance, efficient use of storage, and database activities. For information on SC Series storage and Oracle, see the best practices document, [Dell EMC SC Series Arrays and Oracle](#).

2.2 Oracle technology

ASM is an alternative to other conventional volume managers, file systems, and raw devices. It provides a simplified storage management interface for DBAs. Broadly speaking, ASM is a logical volume manager that virtualizes storage to a set of disk groups and provides redundancy options to enable a high level of protection. Oracle recommends using ASM for storage management.

Flex ASM is an enhanced version of ASM that is offered in Oracle 12c. Flex ASM provides the capability to change the architecture of the ASM cluster, and provides a way to mitigate database instance outages should the ASM instance running on the database server sustain an outage.

With respect to the physical storage, there is no difference on how storage should be configured for ASM or Flex ASM. The configuration can be the same. The difference between ASM and Flex ASM is how ASM is configured. With no difference in physical configuration between ASM or Flex ASM implementations, the use of SC Series storage features remains the same between the ASM and Flex ASM implementations.

2.3 Oracle ASM features

Oracle ASM was introduced in Oracle 10g and provides a vertical integration of the file system and a volume manager specifically built for only Oracle database files. ASM can assemble multiple disk devices that behave as block devices into logical containers called disk groups, and can perform non-intrusive storage configuration, striping, and mirroring within the disk groups with automatic rebalancing. ASM distributes I/O loads across all available disk devices within disk groups to optimize performance and resource utilization, while removing the need for manual I/O tuning by spreading out the database files to avoid hotspots. ASM also helps DBAs manage a dynamic database environment by allowing database growth and transparent data relocation among the disk devices without having to shut down the database to adjust storage allocation.

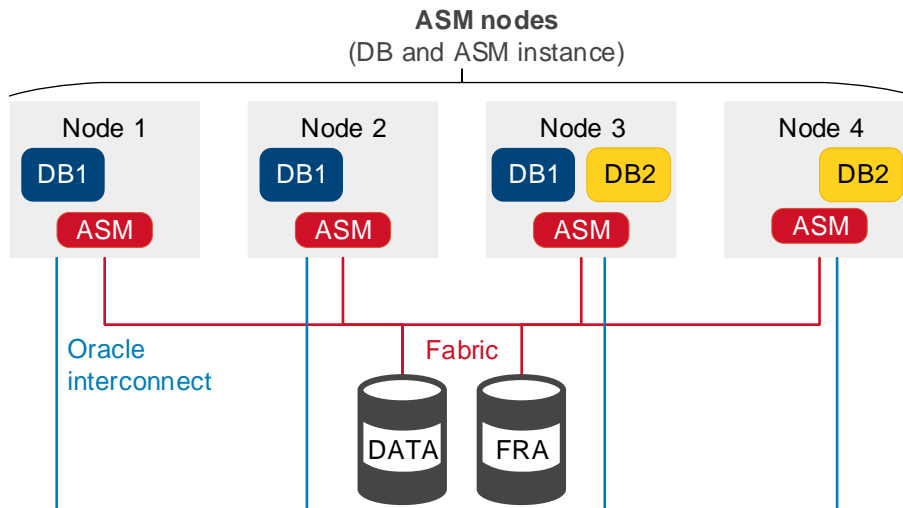


Figure 1 Traditional ASM configuration with dedicated ASM instance for each database server

For LUNs intended for ASM, there are several tools at the OS level that can be used to manage them and provide consistent device naming and permission persistency: ASMLib, ASM Filter Driver (ASMFd), or UDEV. This document briefly reviews ASMLib. For additional information on ASMLib, ASMFd, UDEV, and the configuration and benefits of using ASM, see the document, [Dell EMC SC Series Arrays and Oracle](#).

2.3.1 Oracle Flex ASM features

Prior to the availability of Flex ASM, an ASM instance is required on each database server where there is a database instance. This creates the possibility of a single point of failure with the ASM architecture on any one of the database servers. If an ASM instance fails, so would the corresponding database instance.

Since ASM manages the physical access to ASM files and its metadata, a shutdown of the ASM instance will cause the client database instances to shut down as well. Flex ASM mitigates this issue by starting another ASM instance in the cluster and reroutes ASM traffic to an available ASM instance in the cluster. The selected ASM instance is the one with the least amount of activity. Rerouting ASM traffic resolves the single point of failure in the event that an ASM instance fails, and provides a highly scalable, massively parallel, and highly available framework. ASM traffic can reside on either the private Oracle RAC interconnect, or an optional dedicated network for ASM traffic. The number of nodes running an ASM instance in a Flex ASM configuration is three by default, but is controlled by the ASM cardinality and can be altered.

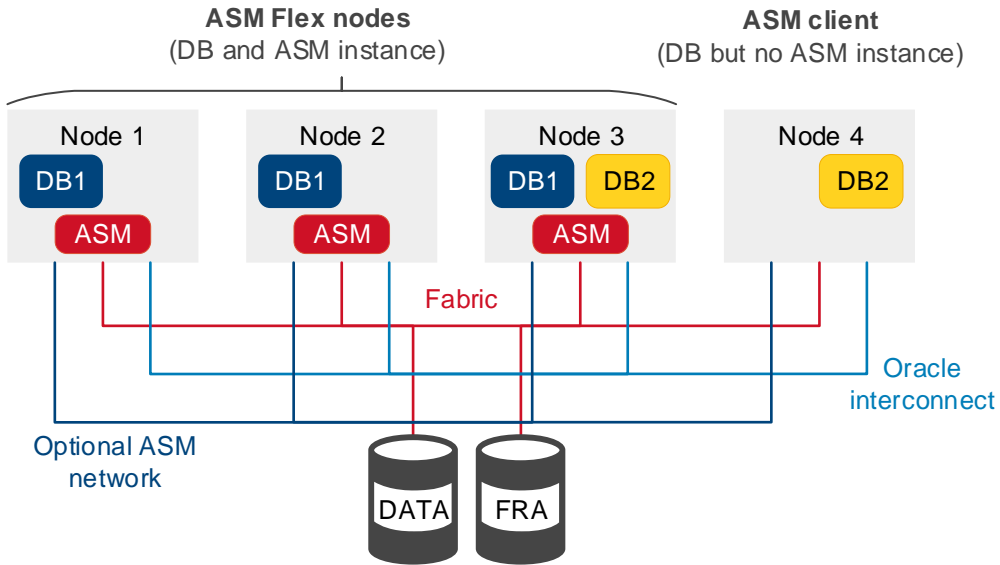


Figure 2 Flex ASM architecture with four database servers with ASM cardinality of three

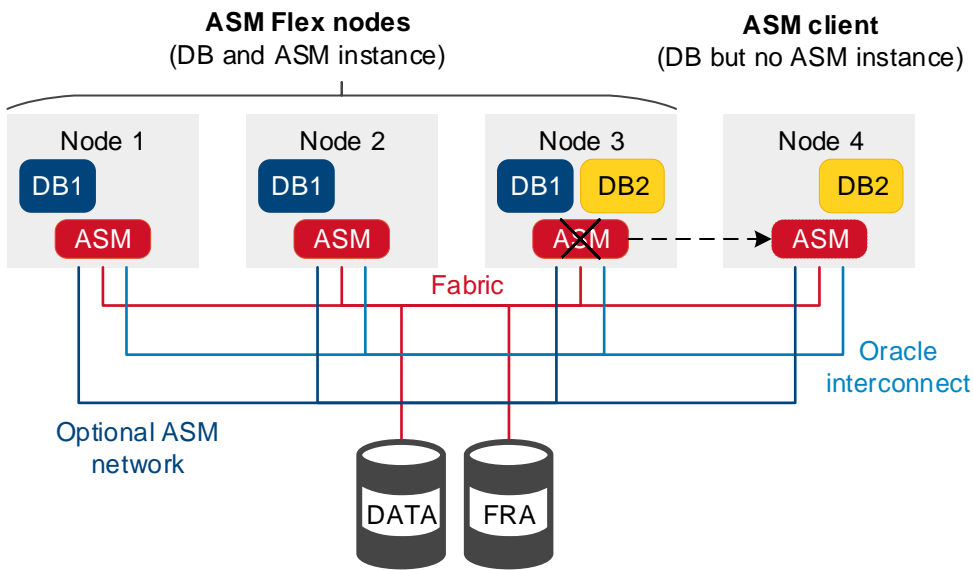


Figure 3 Flex ASM architecture with rerouted ASM traffic

Note: All SC Series volumes intended to be used by ASM must be presented on all ASM Flex nodes and ASM client nodes intended to be used. This applies to a node not currently running a database instance or an ASM instance, but may at some point be used by flex ASM as an ASM instance failover node.

2.3.2 Oracle ASMLib features

Introduced in 2004, ASMLib is an optional set of tools and kernel driver that can be inserted between ASM and the hardware, and contains an application library used by the Oracle database software to access ASM disks in either a standalone or RAC configuration. ASM instances and database instances can use ASMLib as an alternative interface for disk access, yielding greater stability of equivalent workloads on standard Linux file systems. Oracle recommends implementing ASM with ASMLib or ASMFD for better manageability and persistent device naming.

ASMLib has three components:

Kernel driver (oracleasm): Oracle ASMLib Linux kernel driver that is open-sourced (GPL) and is available from Oracle.

Support tools (oracleasm-support): Open-source package (GPL) containing utilities to manage the ASMLib Linux kernel driver and is available from Oracle as an RPM.

Application library (oracleasm package): ASM library. This is a closed source, binary-only RPM, available as a free download from Oracle.

Although it is recommended to use ASMLib or ASMFD, it is not necessary. Should they not be used, udev rules or device mapper multipathing must be used to achieve persistent device naming.

2.3.3 Benefits of ASM, Flex ASM, and ASMLib

Some benefits of using ASM or Flex ASM together with ASMLib are listed in Table 1. For information about additional benefits, refer to the references in appendix D.

Table 1 Benefits of Oracle ASM and ASMLib

Benefit	Details
Simplified RAC implementation	Provides shared storage (raw or file systems)
Recommended by Oracle	ASM is the Oracle-preferred logical volume manager providing greater efficiencies and optimizations when working with database files through direct and async I/O provided by the ASMLib kernel driver
Greater data reliability	ASM provides data mirroring on a file basis rather than on a volume basis, allowing, the same disk group to contain a combination of files protected by mirroring, parity, or not protected at all
Improved performance	Automatic distribution of database files across all disks in a disk group, which delivers the performance of raw disk I/O without the inconvenience of managing raw disks
Reduced overhead	Eliminates the overhead at the file system layer since I/O is passed from the database directly to the disks
Increased database uptimes	Unlike logical volume managers (LVMs), ASM maintenance operations do not require a database outage
Support for multiple stripe sizes	ASM provides the ability to use different stripe sizes for different database file types
Reduced time to resync data after a failure	ASM fast mirror resync reduces the amount of time required to resynchronize a disk after a failure by journaling the ASM extents that are modified during the outage and applying them after the failure is resolved

3 Components of ASM

3.1 ASM instance

Starting with Oracle 10g and above, there are two types of instances: database and ASM. The ASM instance, which is generally named +ASM, is started with the `INSTANCE_TYPE=ASM` `asm init.ora` parameter. This parameter 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 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. For a list and description of the background processes, see the appendices of this document.

3.2 ASMLib

As mentioned in the prior section, Oracle developed two optional storage management API interfaces for block-device management, discovery, and provisioning, one of which is ASMLib. ASMLib provides storage vendors the ability to supply a plug-in library to extend storage-related features and exploit storage array capabilities when using ASM. It is an add-on module that simplifies the management and discovery of ASM disks and provides an alternative to the standard operating system interface for ASM to identify and access block devices.

ASMLib is not required to run ASM, but if ASMLib is not used, either ASMFD or native OS methods (like UDEV rules) must be used to manage the ASM block devices. ASMLib also allows a database to use more efficient access to the logical volume groups it is using.

If ASMLib is used, the user and group permissions will be set correctly on the devices under `/dev/oracleasm/disks/*`. The values of user and group are configured while running `/etc/init.d/oracleasm configure`.

3.3 Disk groups

Disk groups are the primary component of ASM. Once the ASM disk devices are discovered, a disk group can be created that will encapsulate one or more of the ASM disks. Once the disk groups are created, they can be specified as the default location for database files created in the database. Files written on ASM disks are ASM files and can have automatically generated filenames. A disk group is the highest-level data structure in ASM and is comparable to an LVM volume group. There are some differentiators between typical LVM volume groups and ASM disk groups.

- An ASM file system layer is implicitly created within a disk group. This file system is transparent to users and only accessible through ASM, interfacing databases, and the ASM command-line tool.
- There are inherent automatic file-level striping and mirroring capabilities. A database file created within an ASM disk group will have its file extents (not database extents) distributed equally across all online disks in the disk group, which provides an even I/O load.

Oracle recommends that generally no more than two disk groups be maintained and managed per single ASM instance or RAC cluster. However, it is highly recommended to evaluate this recommendation and create additional disk groups as necessary. For a discussion on when additional disk groups might be necessary or desirable, see the discussion on snapshot profiles in the document, [Dell EMC SC Series Arrays and Oracle](#). Data that can be categorized with different storage characteristics can be stored in distinct disk

groups with different redundancy settings and fail-groups if necessary. It is recommended to have at least the following disk groups:

- **Database data (DATA):** Where active database files such as data files, control files, online redo logs, archive redo logs, and change-tracking files used in incremental backups are stored.
- **Flash Recovery Area (FRA):** Where recovery-related files are created, such as multiplexed copies of the current control file and online redo logs, backup sets, and flashback log files.

3.3.1 ASM disk group redundancy

There are three types of redundancy for ASM disk groups:

- **Normal Redundancy:** Provides 2-way mirrored redundancy which requires two ASM failure groups. This is the default.
- **High Redundancy:** Provides 3-way mirrored redundancy which requires three ASM failure groups.
- **External Redundancy:** Provides no protection within ASM. Since ASM does not provide redundancy when external redundancy is used, the server CPU and other resources are less consumed, which means there is an increase of available CPU cycles for database operations. This level of redundancy requires RAID protection on the storage arrays. This is generally the recommended redundancy level when using SC Series storage. There are use cases in Oracle ASM that might require either normal or high redundancy even if SC Series storage is used. SC Series storage provides redundancy at the storage (disks) level. If a database server cannot access the ASM LUN at the operating-system level due to operating-system errors, the ASM instance will fail. If this is a primary concern, configure the ASM disk group with normal redundancy (with two failure groups), or high redundancy. If extended distance clusters are a primary objective, either normal or high redundancy levels must be chosen.

4 ASM best practices with SC Series storage

For best practices with Dell EMC SC Series storage and ASM, see the document, [Dell EMC SC Series Arrays and Oracle](#). This document includes information on:

- Fibre Channel connectivity
- Zoning
- SC Series virtual port mode
- SC Series volume size and quantity
- Dynamic multipathing
- Partitioning or not partitioning ASM disks
- Partition alignment
- Storage profiles
- Snapshot (Replay) profiles

Additional information is provided as follows:

LUN sizes can vary depending on the size of the database. Generally, Dell EMC does not recommend any LUN size greater than 1 TB. If a database is roughly 1 TB in size, 4 x 250 GB LUNs might be a good number of LUNs to start with for the ASM disk groups.

When using Red Hat or Oracle Linux versions 6 or 7, ensure **/etc/multipath.conf** contains a device stanza for Compellent devices and the following default settings:

```
defaults {
    find_multipaths yes
    user_friendly_names yes
    polling_interval 5
}
devices {
    device {
        vendor "COMPELNT"
        product "Compellent Vol"
        features "0"
        hardware_handler "0"
        prio "const"
        no_path_retry 24
        path_grouping_policy "multibus"
        path_checker "tur"
        failback immediate
        rr_weight "uniform"
        fast_io_fail_tmo 5
        dev_loss_tmo infinity
        path_selector "round-robin 0"
    }
}
```

For additional multipath information, see the My Oracle Support (MOS) document 1538626.1, SC Series Storage System Owner's Manual, *Dell Storage Center Deployment Guide*, and [Dell EMC SC Series with Red Hat Enterprise Linux 6x](#) best practices.

If partitioned devices are used, attempting to expand the size of a disk that is not the primary partition requires advanced knowledge of Linux systems and Oracle backups and should only be done after careful planning and consideration. This includes making sure a valid backup exists of the entire database.

When not using ASMLib, it might be advantageous to use an entire disk and not a partition to ease the administration and management of the disk should the disk increase in size. Additional information on resizing a LUN can be found in the [Dell EMC SC Series with Red Hat Enterprise Linux 6x](#) best practices document. If ASMLib is used, the disk must be partitioned per Oracle requirements. The first partition will be used by ASMLib to save its metadata, and the remaining portion of the disk will be used by the database.

Dell EMC recommends creating an even number of LUNs per disk group in a dual-controller array since SC Series storage will automatically balance volumes across both controllers, resulting in I/O activity being more balanced across controllers.

Having fewer disks per ASM disk group minimizes the administration headaches for DBAs, system administrators, and storage administrators. However, consider using more than two disk groups if possible.

If necessary, disk groups may contain more than two ASM disks. Multiple disks provides a way to take advantage of multiple disk queues at the operating-system level and helps if the HBA is not saturated.

If possible, use ASM external redundancy on all ASM disks provisioned from SC Series volumes.

5 Conclusion

In summary, ASM provides a simplified storage management interface for DBAs while SC Series storage technologies provide ease of use for storage provisioning and management for system administrators. These technologies simplify the management of database resources and storage resources. This allows DBAs and system administrators more time to concentrate on other important tasks and reduce the overlap of responsibilities.

A Flex ASM client query

The following 12c query shows what database instances are being serviced by the ASM instance that executed the following query:

```

set linesize 180
set pagesize 60

col group_number          format 99          heading "Grp|Num"
col instance_name        format a8          heading "Inst|Nm"
col db_name              format a8          heading "DB|Name"
col cluster_name         format a12         heading "Clu|Nm"
col status               format a9          heading "Status"
col state                format a7          heading "State"
col name                 format a7          heading "DG Name"
col type                 format a6          heading "ASM|Redun-|dancy"
col total_mb             format 999,999     heading "Total|MB"
col allocation_unit_size format 9,999,999   heading "Alloc|Unit|Sz"
col offline_disks        format 99         heading "Offline|Disks"
col voting_files         format a7          heading "Voting|Files"

select ac.group_number
       , ac.instance_name
       , ac.db_name
       , ac.cluster_name
       , ac.status
       , ac.con_id
       , adg.sector_size
       , adg.block_size
       , adg.allocation_unit_size
       , adg.state
       , adg.type
       , adg.total_mb
       , adg.offline_disks
       , adg.voting_files
from v$asm_diskgroup adg
     , v$asm_client ac
where adg.group_number = ac.group_number
order by 1, 2;

```


The following is sample output of the prior query:

Grp Num	Inst Nm	DB Name	Clu Nm	Status	Alloc			ASM		
					Unit Sz	State	Redun- dancy	Total MB	Offline Disks	Voting Files
1	+ASM1	+ASM	rac-cluster	CONNECTED	1,048,576	MOUNTED	EXTERN	15,358	0	Y
1	-MGMTDB	_mgmtadb	rac-cluster	CONNECTED	1,048,576	MOUNTED	EXTERN	15,358	0	Y
2	+ASM1	+ASM	rac-cluster	CONNECTED	1,048,576	MOUNTED	EXTERN	511,993	0	N
2	prd1	prd	rac-cluster	CONNECTED	1,048,576	MOUNTED	EXTERN	511,993	0	N
2	prd2	prd	rac-cluster	CONNECTED	1,048,576	MOUNTED	EXTERN	511,993	0	N
3	prd1	prd	rac-cluster	CONNECTED	1,048,576	MOUNTED	EXTERN	511,993	0	N
3	prd2	prd	rac-cluster	CONNECTED	1,048,576	MOUNTED	EXTERN	511,993	0	N

7 rows selected.

The output shows that instance +ASM1 is servicing three database instances:

- -MGMTDB
- prd1
- prd2

However, a quick check for pmon processes on the server shows that database instances +ASM1 and prd1 are the only instances running on the server.

```
ps -ef | grep pmon | grep -v grep
grid      102713      1  0 10:36 ?          00:00:00 asm_pmon_+ASM1
oracle    103746      1  0 10:37 ?          00:00:00 ora_pmon_prd1
```

Instances prd2 and -MGMTDB are running on a different server.

B Other useful commands for an ASM environment

The following commands can be useful in an ASM configuration. For a complete list of commands, refer to the references in appendix D.

Component	Description
srvctl modify asm -count <n>	Change cardinality of Flex ASM instances to <n>
srvctl config asm	Display ASM configuration
srvctl status asm -detail	Display status of ASM on all nodes
crsctl status res ora.asm [-t -l -f]	Display status of ASM and cardinality of Flex ASM if enabled.
crsctl status res ora.ASMNET1LSNR_ASM.lsnr [-t -l -f]	Display status of the listener used by Flex ASM
asmcmd showversion	Display ASM version
asmcmd showpatches	Display the patches applied on the grid home
asmcmd showclustermode	Show if flex ASM is enabled
asmcmd showclusterstate	Displays ASM rolling migration/patching state
asmcmd lsct <diskgroup>	Display information about ASM clients
asmcmd pwget --asm	Display ASM password file location
asmcmd spget	Display ASM spfile file location
asmcmd shutdown	Shut down an Oracle ASM instance
asmcmd startup	Start up an Oracle ASM instance
asmcmd spbackup	Back up an Oracle ASM spfile
oifcfg getif	Display the network interfaces (including ASM's)
olsnodes -n -s -a	Display state and type of nodes (hub,leaf) in cluster

See the *Oracle Automatic Storage Management Administrator's Guide*, 12c Release 1 (12.1), for additional information.

C 512 byte and 4K byte sectors

Since consumer-grade HDDs were first introduced in the early 1980s, 512 byte sectors were used. But as areal density requirements kept increasing on disks, there was a need to increase the sector size. By 2011, all hard drive manufacturers committed to ship new hard drive platforms that supported a new format called Advanced Format (AF).

Advanced Format uses storage surfaces more efficiently for large files, and allows for the integration of stronger error correction algorithms to maintain data integrity at higher storage densities. These disks would be presented to the OS with physical block size of 4K bytes.

Up to a point, Oracle only supported devices presented to the OS with 512 byte physical sectors. With the advent of the 4K byte physical sectors, Oracle binaries had to change. Depending on the Oracle configuration, including versions of Oracle software, there are a couple of issues using 4K byte physical sectors with Oracle and a couple of options available to mitigate the issues. With respect to the 12c (12.1.0.2.0) environment used for this paper, the only issue discovered pertained to writing spfiles to disk devices that are presented to the OS with 4K byte physical sectors. This caused issues for the Grid installer when it tried to create its management repository database, and during the creation of a database using DBCA.

To mitigate any issue caused by 4K or 512 byte physical sectors with ASMLib, Dell EMC recommends that the following options be executed in order listed and that Oracle Support is contacted for assistance.

1. If ASMLib is used, it might be possible to configure `/etc/sysconfig/oracleasm` to direct the ASMLib driver to use the size of the logical block rather than the size of the physical block. The ASMLib directive is not supported in all version of Oracle, so check My Oracle Support if it can be used.

```
ORACLEASM_USE_LOGICAL_BLOCK_SIZE=true
```

2. Oracle patches are available for download. For a complete list of issues and available patches pertaining to 4K byte sectors, search for **4K vs 512 byte sectors** in My Oracle Support. Some of the notes will make reference to specific patches that must be applied to the Oracle environment to use 4K byte sectors.
3. Consider using UDEV rather than ASMLib.
4. See the My Oracle Support notes and threads that reference a workaround which makes sure the storage provider would present the logical and physical block size of volumes that reside on Advanced Format devices to be 512 bytes. If Oracle Support is unable to resolve issues related to using 4K byte blocks, contact Dell Copilot for assistance.

D Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware, and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC storage platforms.

D.1 Related resources

See the following list of referenced or recommended Dell EMC publications:

- [Dell EMC SC Series with Red Hat Enterprise Linux 6x](#)
- [Dell EMC SC Series Arrays and Oracle](#)
- [How to deploy Oracle 11gR2 on RHEL6/Oracle Linux 6](#)

The following include referenced or recommended Oracle publications:

- [Oracle Linux Installation Guide for Release 6](#)
- [Oracle Linux Administrator's Guide for Release 6](#)
- [Oracle Grid Infrastructure, Installation Guide, 12c Release 1 \(12.1\) for Linux, E48914-16](#)
- [ASM Enhancements in Oracle Database 12c](#)
- [Oracle Database 11gR2 Database Administration Database Performance Tuning Guide](#)
- [Oracle Database Storage Administrators Guide](#)
- [Oracle 11g Database Concepts](#)
- [Oracle Database Release Notes 11gR2 for Linux](#)
- [Oracle Database Installation Guide 11gR2 oc.htm](#)
- [Oracle Grid Infrastructure Installation Guide 11gR2](#)
- [Oracle Automatic Storage Management, Administrator's Guide, 12c Release 1 \(12.1\)](#)
- Grid installation reports error on "Prerequisite checks" PRVF -5150 [ID 1498115.1]
- Device Checks for ASM Fails with PRVF-5150: Path ORCL: is not a valid path [ID 1210863.1]
- Bug 14048512 - PRVF-5150 : Path ORCL:DISK1 is not a valid path on all nodes [ID 14048512.8]
- GI 11.2.0.3 Installation On Errors PRVF-5150 saying ASMLIB path is not a valid path [ID 1474961.1]