

An Oracle White Paper  
October 2011 Update

# Oracle Database 11g Release 2 Extending Automatic Storage Management to Manage All Data

Executive Overview .....	1
Introduction .....	1
One Integrated Solution for ALL Data.....	3
ASM Cluster File System (ACFS).....	4
ASM Dynamic Volume Manager.....	5
ACFS Snapshots.....	6
Oracle Clusterware OCR and Voting File Support.....	7
Extensions for System and Database Administrators .....	7
ASM Installation and Configuration Tools .....	7
Full Features ASMCMD Extensions.....	8
Functional Enhancements .....	9
ASM File Access Control (ACL).....	9
ASM Disk Group RENAME.....	9
ASM File Mapping .....	10
ASM 4k Sector Size Support .....	10
Tunable Performance With Oracle Disk Placement (ODP) .....	11

## Executive Overview

As storage requirements double every 18 months, Oracle customers continue to deal with complex storage management challenges in their data centers. Convention methods to manage storage are too complex and costly because the customer is typically forced to be the integrator for multiple tools and software layers when it pertains to choosing their storage hardware platforms as well as volume and file system management and clustering platforms. This problem is compounded when it comes to effectively managing their structured and non-structured data on multiple platforms having to deal with multiple vendors for support, learning different tools with varying user interfaces and procedures for installation, configuration, performance tuning and management in general.

## Introduction

Automatic Storage Management (ASM) feature in Oracle Database 11g Release 2 extends ASM functionality to manage ALL data: Oracle database files, Oracle Clusterware files and non-structured general purpose data such as binaries, externals files and text files. ASM simplifies, automates and reduces cost and overhead by providing a unified and integrated solution stack for all your file management needs eliminating the need for 3<sup>rd</sup> party volume managers, file systems and clusterware platforms.

ASM 11g Release 2 introduces key enhancements and extensions that further simplify storage management and provide value for Oracle customers.

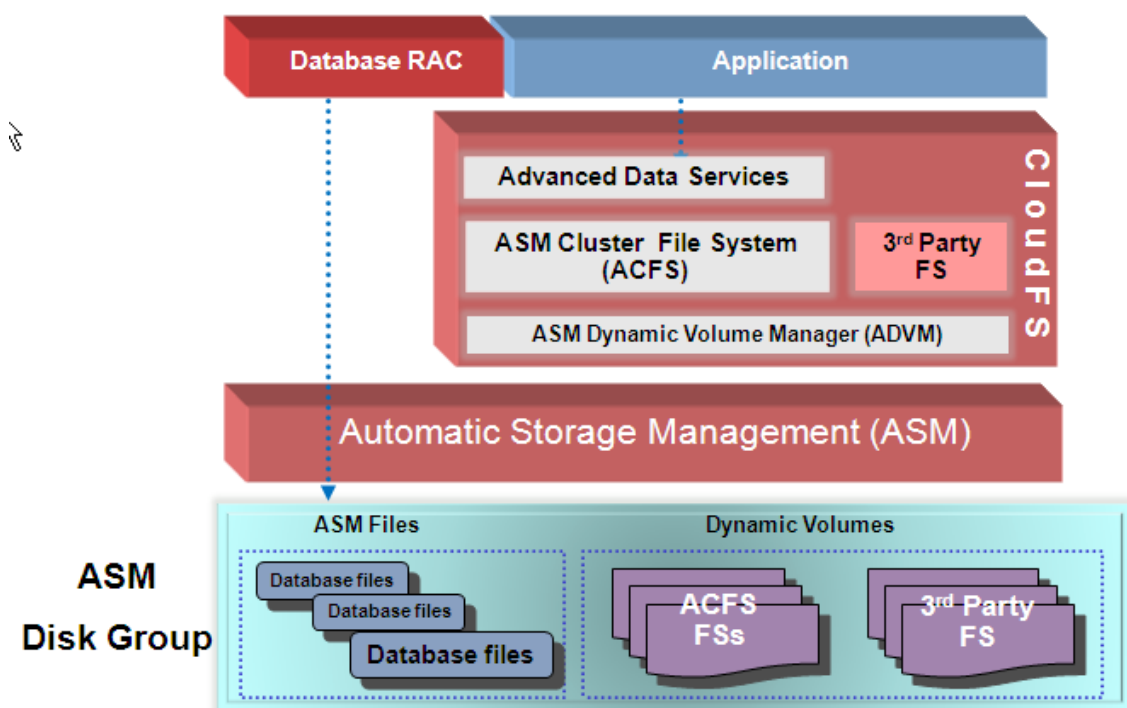
- General purpose ASM Cluster File System (ACFS)
- ASM Dynamic Volume Manager (ADVM)
- Snapshot data service for ACFS
- Oracle Clusterware OCR and Voting File Support
- Extensions for system and database administrators
- Functional Enhancements
- Tunable performance

Diagram 1 illustrates the high level ASM architecture as an integrated solution that fulfills the requirements for Oracle databases and Real Application Clusters natively, and provides a platform for all other applications and 3<sup>rd</sup> party file systems as well.

The ASM instance continues to be the most efficient way to manage Oracle database files directly. The ASM instance manages disk groups (pools of storage) with highest automation, ease and performance. ASM provides a consolidated environment where multiple databases can be stored and managed within the same set of disk groups. In addition, the Oracle Cluster Registry (OCR) and Voting files can now be automatically created by ASM and managed with high integrity and availability.

ASM also introduces a dynamic loadable kernel volume manager called the ASM Dynamic Volume Manager (ADVM) that provides a general purpose volume management platform not only for ACFS file systems but also for 3<sup>rd</sup> party file systems such ext3 for Linux. ASM Dynamic Volumes leverage all the powerful ASM features such as storage provisioning, rebalancing, redundancy and automation and are managed by the ASM instance.

The ASM Cluster File System is a new POSIX/X-OPEN compliant general-purpose scalable storage management technology that extends the ASM functionality to support ALL non-Oracle database files for Linux and Windows platforms. Examples of files supported by ACFS are Oracle binaries, application executables, trace files, alert logs, BFILEs, audio/video/image file and any other general-purpose files. ACFS Snapshot is the 1<sup>st</sup> value added service being introduced in this release. It is a read-only space efficient point-in-time snapshot technology for ACFS file systems.



## One Integrated Solution for ALL Data

Oracle customers can now benefit from a single integrated solution instead of acting as the systems integrator and dealing with the complexities of integrating different software layers from different vendors in house. ASM/ADVM/ACFS/Snapshot is tightly integrated with the Oracle Database, RAC, Clusterware and other Oracle technologies and presents a comprehensive and well tested solution to our customers.

ASM feature in Oracle Database 11g Release 2 is available across all Oracle supported OS platforms (ACFS /ADVM initially available on Linux and Windows) and managed with the same user interface across platforms. ASM benefits from one management interface, one installation and configuration tools, one Clusterware framework and one vendor for support.

ASM is designed to be always-on-line increasing database availability despite storage and ASM configuration changes. ASM continues to bring innovations in automation to further simplify file and storage management for an Oracle database server environment. The ASM architecture provides for highest performance and balanced I/O distribution on a continual basis and without tedious complex administration. It is a scalable solution that makes it suitable for deploying very small as well as multi-hundred terabyte databases. ASM disk groups function as storage containers that are ideal for consolidation of databases and file systems that result in optimal storage utilization reducing the storage waste in your computing environments.

## ASM Cluster File System (ACFS)

ASM Cluster File System (ACFS) is a general-purpose, single node and cluster-wide file system for non-Oracle database files. ACFS is POSIX and X/OPEN compliant file system for Linux and Unix servers and as a windows file system for Windows systems. This means that ACFS is managed using native operating application programming interfaces (API) and command line (CLI) tools. An ACFS file system is created on top of an ASM Dynamic Volume (covered in the next section). Dynamic Volumes are simply ASM files (with an ASM volume type) and inherit and benefit from all ASM functionality as an integrated solution.

ACFS supports large files with 64 bit file and file system data structure leading to exabyte file system capacities on 64 bit platforms. Variable extent based storage allocation, direct path file I/O, and high performance directory architecture contributes to fast performance. File system integrity and fast recovery is achieved through metadata checksum and journaling. ACFS is designed as a peer to peer, multi-node, shared file system model and delivers coherent, cache, direct path file I/O to ASM storage from each cluster member.

ACFS leverages core ASM features and functionality that enables dynamic file system resizing, maximized performance through direct access to ASM disks, balanced distribution of ACFS files across ASM disks for even I/O distribution, and data redundancy through ASM mirroring protection mechanism.

ACFS can be managed using command-line tools (OS native, ASMCMD and ACFSUTIL), Oracle Enterprise Manager, ASM Configuration Assistant (ASMCA) and SQL command interface. ACFS can also be configured for access by remote clients using industry standard NAS file access protocols: NFS and CIFS. ACFS provides for single-node Oracle restart mount registry (similar to fstab in Linux) and the Oracle Grid Infrastructure cluster-wide mount registry, dynamic file system resizing, and multiple space-efficient snapshots for each file system. An ACFS file system is automatically mounted if inserted into the mount registry.

Simple steps to create, provision capacity and manage an ACFS file system:

Create a 10 GB ASM volume and a file system and register on ACFS mount registry:

```
$ ASMCMD volcreate -d DATA -s 10G volume1 ; Create volume 'volume1'  
# mkfs -t acfs -b 4k /dev/asm/volume1 ; Create acfs file system
```

```
# mount -t acfs /dev/asm/volume1 /oracle/cluster1/myacfs ; Mount file system  
# acfsutil registry -a /dev/asm/volume1 /oracle/cluster1/myacfs ; Write to registry
```

Provision capacity to expand volume1 and myacfs by 5GB:

```
$ ASMCMD volresize DATA volume1 -s 5G
```

ASM leverages its rebalancing feature to automatically distribute all ACFS data evenly across the expanded storage capacity ensuring even distribution of I/O and high performance.

## ASM Dynamic Volume Manager

The ASM Dynamic Volume Manager (ADVM) provides volume management services and a standard device driver interface to its clients (ACFS, ext3, OCFS2...). An ASM Dynamic Volume is constructed out of an ASM file with an 'asmvol' type attribute to differentiate it from other ASM file types (i.e. datafile, redolog, etc...). ACFS and 3<sup>rd</sup> party file systems can use ADVM as a volume management platform to create and manage file systems leveraging all the power and functionality of ASM features. Any number of Dynamic Volumes can be created in an ASM disk group.

```
$ ASMCMD volcreate -d DATADG -s 10G volume1
```

An ASM volume device is created on the OS automatically when an ASM Dynamic Volume is created in a disk group using either ASMCMD, SQL or the Enterprise Manager. An ADVM volume device exports ASM volume management features and ensuring dynamic volume integrity in the event of abnormal failures.

ASM device file name examples are as follows:

Linux: /dev/asm/DATADG/volume1

Windows: asm\_DATADG\_volume1

ADVM implements its own extent and striping algorithm to ensure highest performance for general purpose files. An ADVM volume is four columns of 64MB extents and 128KB stripe width by default. ADVM writes data in 128KB stripes in a round robin fashion to each column before starting on the next four column extents. ADVM uses Dirty Region Logging (DRL) for mirror recovery after a node or possibly an instance failure. This DRL scheme requires a DRL file in the ASM disk group to be associated with each ASM

Dynamic Volume file. ADVN functionality includes: Create, resize, delete, enable, disable, set and info.

## ACFS Snapshots

Automatic Storage Management Cluster File System read-write Snapshots may be used in test and development and recover from inadvertent modification or deletion of files from a file system. A snapshot can also be used as the point-in-time source for a file system backup, as it can be created on demand to deliver a current, consistent, on-line view of an active file system.

An Automatic Storage Management Cluster File System Snapshot is a read-write on-line, space efficient, point in time copy of the file system. The snapshot copy is initially sparse and merely references the storage allocation information maintained by the source file system and therefore, requires minimal storage capacity because it maintains only the changed blocks. An Automatic Storage Management Cluster File System Snapshot utilizes a Copy-On-Write (COW) technology to maintain point in time view of the file system by copying the current extent value to the snapshot before modifying or deleting the source file extent.

ACFS snapshots are immediately available for use after they are created. Up to 63 snapshot views are supported with this release. Snapshots are always on-line while the file system is mounted. ACFS snapshots can be used to recover from inadvertent modification or deletion of files from a file system. A snapshot can also be used as the source for a file system backup, as it can be created on demand to deliver a current, consistent, on-line view of an active file system.

ACFS snapshots are created in the .ACFS/snap directory automatically which simplifies and eliminates the need for maintaining a separate storage pool for the snapshots. ACFS snapshots are administered with the 'acfsutil snap' commands as well as the Enterprise Manager database/grid control.

```
$ acfsutil snap create snap_name mount_point
```

A new V\$ASM\_ACFSSNAPSHOTS view is also available to query purposes.



## Oracle Clusterware OCR and Voting File Support

Starting with ASM 11g Release 2, the Oracle cluster Registry (OCR), Voting file and ASM SPFILE can be stored in an ASM disk group. ASM simplifies and automates the creation of OCR/Voting files and eliminates the need to create separate partitions or use expansive cluster volume managers and file systems for that purpose.

The OCR configuration file is treated as a regular ASM file. Therefore, it benefits from all ASM file functionality and attributes such as redundancy and re-balancing. The OCR redundancy follows the disk group mirroring choices.

The Voting file is automatically created on exclusive regions of ASM disks by default. ASM creates 1/3/5 Voting files based on respective External/Normal/High disk group redundancy selected. A Quorum failure group is optionally used to store a voting file exclusively (a quorum FG can be used for database files). The Quorum failure group may be created in configurations where only two failure groups exist for a normal/high redundancy disk group or the user chooses to segregate the Voting file from the database files.

In the event that an ASM disk containing a Voting file fails, ASM will automatically re-create the Voting file on a different ASM disk to maintain redundancy. The OCR file recovers leveraging the ASM disk group mirroring technology.

The crsctl remain the tools that can be used to manage OCR and Voting files.

## Extensions for System and Database Administrators

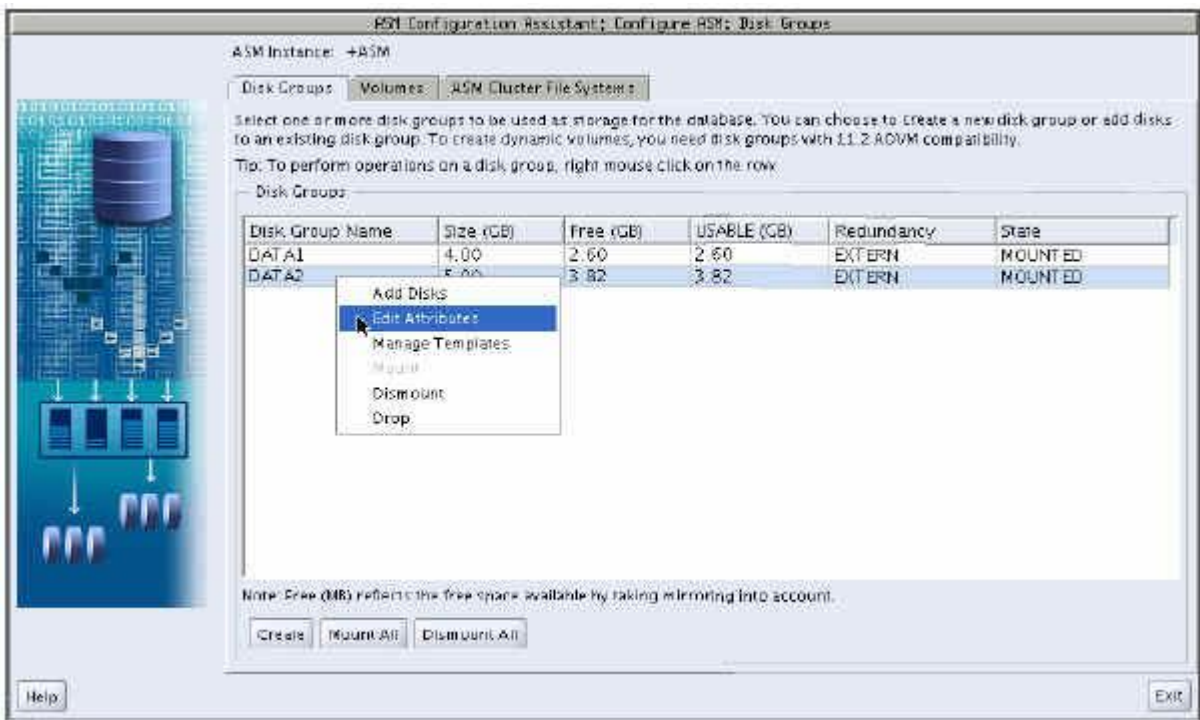
### ASM Installation and Configuration Tools

New installation and configuration tools have been developed to allow ASM and Oracle Clusterware installation and configuration independent of Oracle databases. This provides the flexibility for aligning roles and responsibilities for organizations where system administrators are tasked with managing infrastructure software such as ASM and Clusterware while DBAs continue to have responsibility for Oracle Database installation and configurations.

Oracle Universal Installer (OUI) and ASM Configuration Assistant (ASMCA) are new system admin friendly tools and are used to install ASM and Oracle Clusterware in a

single Grid Infrastructure home 'Grid'. ASM is no longer a part of the database install. This OUI installation has the option to minimally configure ASM and ACFS in its process. This also automates creation of an Oracle Database home on ACFS in preparation for a database home to be installed in the database installation phase.

ASMCA is a new graphical tool that is designed to simplify ASM configuration management especially for system administrators. ASMCA simplifies initial configuration and management for ASM disk groups, ASM Dynamic volumes and ACFS file systems. ASM Disk groups, volumes and file systems can be created and managed efficiently.



### Full Features ASMCMD Extensions

The popular ASM command line interface ASMCMD is enhanced to provide full management functionality paralleling the SQL\*Plus commands. This is a significant step for providing the agility needed for system administrators and DBAs to manage the entire ASM environment with their favorite management API. ASMCMD now provides management commands in the following categories:

- ASM Instance Management (startup, shutdown, ...)
- ASM Disk Group, Disk and Failure Group Management (create, mount, add, drop, ...)
- ASM File Management
- ASMCMD Volume Management (create, resize, ...)
- ASM File Access Control
- Template Management (list, add alter drop, ...)
- User/Group Management (add user, change password, ...)
- ASM Attributes (list, set)
- Miscellaneous (loststat, df, lsof, lsod)

## Functional Enhancements

### ASM File Access Control (ACL)

ASM File Access Control restricts the access of files to specific ASM clients that connect as SYSDBA. An ASM client is typically a database, which is identified as the user that owns the database instance home. ASM File Access Control uses this user name to identify a database. ASM introduces three classes of permissions: owner, group, other and three levels of permissions per class: none, read-only, read-write. You should set the disk group attributes before creating any files in a new disk group. Otherwise, you must explicitly set the permissions and ownership on existing files in a disk group.

To manage ASM File Access Control for a disk group, you need to set the `ACCESS_CONTROL.ENABLED` and `ACCESS_CONTROL.UMASK` disk group attributes when altering the disk group. The `.umask` attribute determines which permissions are masked out on the creation of an ASM file for the user that owns the file, users in the same user group, and others not in the user group. This attribute applies to all files on a disk group. The values can be combinations of three digits `{0|2|6} {0|2|6} {0|2|6}`. The default is `066`.

### Examples

The following are examples of `chmod` and `chown` commands that change permissions and ownership of specific files:

```
ASMCMD> chmod ug+rw +data/orcl/controlfile/Current.260.684924747
```

```
ASMCMD> chown oracle1:asm_users +data/orcl/controlfile/Current.260.684924747
```

### ASM Disk Group RENAME

It is necessary to change the name of an ASM disk group after cloning if it is going to be mounted on the same server for backup, test or other application reasons. The disk group should be dismounted before the rename operation. The *renamedg* tool optionally renames a disk group in two phases:

Phase 1: This step generates a configuration file. Each line contains description of change to be made to one disk.

Phase 2: This step uses the configuration file to perform the renaming of the disk group.

```
$ renamedg -phase one -dgname data2 -newdgname data12 -config /tmp/data2.conf  
$ renamedg -phase two -config /tmp/data2.conf
```

### ASM File Mapping

Oracle Database provides a mechanism to show a complete mapping of a database file to intermediate layers of logical volumes, file systems to actual physical devices. This is accomplished through a set of dynamic performance views (V\$ views). Using these views, one can locate the exact disk on which any block of a file resides.

This mapping is accomplished using a set of internal (FMON) and external (FMPUTL) processes and vendor provided mapping libraries. FMON is responsible for building the mapping information stored in the SGA and includes the following structures: Files, File system extents, Elements, Sub-elements.

ASM provides an efficient Oracle File Mapping interface. For an ASM file name like '+dg/f.1', FMON can directly call ASM mapping functions that are implemented inside ASM to obtain mapping information for ASM disk groups and disks. However, FMON still needs to communicate with FMPUTL for mapping lower level structures provided by vendor libraries.

### ASM 4k Sector Size Support

ASM provides support for 4k sector disk drives without negatively affecting performance. The SECTOR\_SIZE disk group attribute can only be set during disk group creation. The value of SECTOR\_SIZE can be 512, 4096, or 4k. The default value is platform dependent. ASM performs validations to ensure the same disk sector size in a disk group. The

compatible .asm and .rdbms must be advanced to 11.2 for this feature. The following is an example of creating a disk group with 4k sector size:

```
CREATE DISKGROUP data NORMAL REDUNDANCY
FAILGROUP Ctrl_a DISK
'/devices/diska1', '/devices/diska2'
FAILGROUP ctrl_b DISK
'/devices/diskb1', '/devices/diskb2'
ATTRIBUTE 'compatible.asm' = '11.2', 'compatible.rdbms' = '11.2',
          "sector_size"='4096';
```

### Tunable Performance With Intelligent Data Placement (IDP)

It is a well known fact that the outer regions of disk drive platters have much better performance than the inner regions. ASM Optimal Disk Placement (ODP) enables you to specify disk regions on ASM disks for best performance. Using the disk region settings, you can ensure that frequently accessed data is placed on the outermost tracks (hot region) that have greater speed and higher bandwidth. In addition, files with similar access patterns are located physically close together, reducing latency. ODP also allows the placement of primary and mirror extents into different hot or cold regions.

ASM Optimal Disk Placement settings can be specified for a file or in disk group templates. The disk region settings can be modified after the disk group has been created. The disk region setting can improve I/O performance by placing more frequently accessed data in regions furthest from the spindle, while reducing your cost by increasing the usable space on a disk.

## Conclusion

Oracle is offering a comprehensive file management solution with ASM 11g Release 2 to further simplify storage management. ASM is enhanced and optimized to manage structured (database files) as well as non-structured (general purpose files) data within the same framework. The ASM general purpose file system (ACFS) will initially support Linux and Windows with a roadmap to support all OS platforms that the database supports. ASM management is streamlined through one management interface, one install and configuration tool, one cluster framework and one vendor for support. ASM is a scalable solution designed to reduce the total cost of ownership, simplify management,

provide high performance without the need for tuning, increase storage utilization through consolidation and increase database and file system availability.



Extending Automatic Storage Management to  
Manage ALL Data  
October 2011 Update  
Author: Ara Shakian  
Contributing Authors:

Oracle Corporation  
World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200  
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2009, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.