

Oracle® Database
Storage Administrator's Guide
11g Release 2 (11.2)
E10500-01

August 2009

Oracle Database Storage Administrator's Guide, 11g Release 2 (11.2)

E10500-01

Copyright © 2007, 2009, Oracle and/or its affiliates. All rights reserved.

Contributing Authors: Jerry Abraham, Prasad Bagal, Mark Bauer, Eric Belden, Bill Bridge, Chan Chan, Xiaobin Chen, Hanlin Chien, Jonathan Creighton, Steve Fogel, Dave Friedman, Barb Glover, Fred Glover, Shie-rei Huang, Diane Lebel, Rich Long, Colin McGregor, Gregory Mengel, Valarie Moore, Darshan Nagarajappa, Harish Nandyala, Bob Nelson, Satish Panchumarthy, Kathy Rich, Francisco Sanchez, Santhosh Selvaraj, Ara Shakian, Duane Smith, Joshua Smith, Malai Stalin, Amar Subba, Yaser Suleiman, Michael Timpanaro-Perrotta, Nitin Vengurlekar, Anton Verhulst, Radek Vingralek, Ron Weiss, Rajiv Wickremesinghe, Jim A. Williams, Sivakumar Yarlagadda, Song Ye, Krishnan Yegnashankaran, Hector Yuen

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

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

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	xxi
Audience	xxi
Documentation Accessibility	xxi
Related Documents	xxii
Conventions	xxii
What's New in Oracle Database Storage?	xxiii
New Storage Features for Oracle Database 11g Release 2 (11.2)	xxiii
New Storage Features for Oracle Database 11g Release 1 (11.1)	xxvi
1 Introduction to Oracle Automatic Storage Management	
Overview of Oracle Automatic Storage Management	1-1
Understanding Oracle ASM Concepts	1-2
About Oracle ASM Instances	1-2
About Oracle ASM Disk Groups	1-6
About Mirroring and Failure Groups	1-6
About Oracle ASM Disks	1-7
Allocation Units	1-7
About Oracle ASM Files	1-7
Extents	1-8
Oracle ASM Striping	1-9
File Templates	1-9
Understanding Oracle ASM Disk Group Administration	1-9
About Discovering Disks	1-10
About Mounting Disk Groups	1-10
About Adding and Dropping Disks	1-10
About Online Storage Reconfigurations and Dynamic Rebalancing	1-10
2 Considerations for Oracle ASM Storage	
Storage Resources for Disk Groups	2-1
Oracle ASM and Multipathing	2-3
Recommendations for Storage Preparation	2-4

3 Administering Oracle ASM Instances

Operating With Different Releases of Oracle ASM and Database Instances Simultaneously	3-1
Configuring Initialization Parameters for an Oracle ASM Instance	3-2
Initialization Parameter Files for an Oracle ASM Instance	3-3
Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File	3-4
Setting Oracle ASM Initialization Parameters	3-5
Automatic Memory Management for Oracle ASM	3-5
Oracle ASM Parameter Setting Recommendations	3-6
ASM_DISKGROUPS	3-7
ASM_DISKSTRING	3-7
ASM_POWER_LIMIT	3-8
ASM_PREFERRED_READ_FAILURE_GROUPS	3-8
DB_CACHE_SIZE	3-9
DIAGNOSTIC_DEST	3-9
INSTANCE_TYPE	3-9
LARGE_POOL_SIZE	3-10
PROCESSES	3-10
REMOTE_LOGIN_PASSWORDFILE	3-10
SHARED_POOL_SIZE	3-10
Setting Database Initialization Parameters for Use with Oracle ASM	3-11
Managing Oracle ASM Instances	3-12
Administering Oracle ASM Instances with Server Control Utility	3-12
Using Oracle Restart	3-12
Starting Up an Oracle ASM Instance	3-13
About Mounting Disk Groups at Startup	3-15
About Restricted Mode	3-15
Shutting Down an Oracle ASM Instance	3-16
Upgrading an Oracle ASM Instance With Oracle Universal Installer	3-17
Downgrading an Oracle ASM Instance	3-19
Active Session History Sampling for Oracle ASM	3-21
Using Oracle ASM Rolling Upgrade	3-21
Patching Oracle ASM Instances	3-22
Authentication for Accessing Oracle ASM Instances	3-22
About Privileges for Oracle ASM	3-23
Using One Operating System Group for Oracle ASM Users	3-23
Using Separate Operating System Groups for Oracle ASM Users	3-24
The SYSASM Privilege for Administering Oracle ASM	3-25
The SYSDBA Privilege for Managing Oracle ASM Components	3-26
Creating Users with the SYSASM Privilege	3-26
Operating System Authentication for Oracle ASM	3-26
Password File Authentication for Oracle ASM	3-27
Migrating a Database to Use Oracle ASM	3-27
Using Oracle Enterprise Manager to Migrate Databases to Oracle ASM	3-28
Using Oracle Recovery Manager to Migrate Databases to Oracle ASM	3-28
Best Practices White Papers on Migrating to Oracle ASM	3-28

4 Administering Oracle ASM Disk Groups

Disk Group Attributes	4-2
Creating Disk Groups	4-3
Using the CREATE DISKGROUP SQL Statement	4-3
Example: Creating a Disk Group	4-5
Creating Disk Groups for a New Oracle Installation.....	4-6
Specifying the Allocation Unit Size	4-7
Specifying the Sector Size for Drives.....	4-8
Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups.....	4-10
Altering Disk Groups	4-11
Adding Disks to a Disk Group.....	4-12
Example: Adding Disks to a Disk Group	4-12
Adding Volumes to Disk Groups	4-13
Dropping Disks from Disk Groups	4-14
Example: Dropping Disks from Disk Groups.....	4-15
Intelligent Data Placement.....	4-15
Resizing Disks in Disk Groups.....	4-16
Example: Resizing Disks in Disk Groups	4-16
Undropping Disks in Disk Groups.....	4-17
Example: Undropping Disks in Disk Groups	4-17
Manually Rebalancing Disk Groups	4-17
Example: Manually Rebalancing a Disk Group	4-18
Tuning Rebalance Operations	4-18
Oracle ASM Disk Discovery	4-18
How A Disk is Discovered.....	4-19
Disk Discovery Rules.....	4-20
Improving Disk Discovery Time.....	4-20
Managing Capacity in Disk Groups	4-21
Negative Values of USABLE_FILE_MB.....	4-23
Oracle ASM Mirroring and Disk Group Redundancy	4-23
Oracle ASM Mirroring and Failure Groups	4-23
Oracle ASM Failure Groups	4-24
How Oracle ASM Manages Disk Failures.....	4-25
Guidelines for Using Failure Groups.....	4-25
Failure Group Frequently Asked Questions.....	4-25
Oracle ASM Recovery from Read and Write I/O Errors	4-26
Oracle ASM Fast Mirror Resync	4-27
Preferred Read Failure Groups	4-29
Configuring and Administering Preferred Read Failure Groups	4-29
Performance and Scalability Considerations for Disk Groups	4-30
Determining the Number of Disk Groups.....	4-31
Performance Characteristics When Grouping Disks	4-31
Oracle ASM Storage Limits	4-31
Disk Group Compatibility	4-31
Overview of Disk Group Compatibility	4-32
Disk Group Compatibility Attributes	4-33
COMPATIBLE.ASM	4-33

COMPATIBLE.RDBMS	4-33
COMPATIBLE.ADVM	4-33
Setting Disk Group Compatibility Attributes	4-34
Valid Combinations of Compatibility Attribute Settings	4-34
Using CREATE DISKGROUP with Compatibility Attributes	4-35
Using ALTER DISKGROUP with Compatibility Attributes	4-35
Viewing Compatibility Attribute Settings	4-36
Features Enabled By Disk Group Compatibility Attribute Settings	4-36
Reverting Disk Group Compatibility	4-37
Considerations When Setting Disk Group Compatibility in Replicated Environments	4-38
Managing Oracle ASM File Access Control for Disk Groups	4-38
About Oracle ASM File Access Control.....	4-39
Using SQL Statements to Set Disk Group Attributes for Oracle ASM File Access Control.	4-40
Using SQL Statements to Manage Oracle ASM File Access Control.....	4-41
Mounting and Dismounting Disk Groups	4-42
Mounting Disk Groups Using the FORCE Option.....	4-43
Checking the Internal Consistency of Disk Group Metadata	4-44
Dropping Disk Groups	4-44
Renaming Disks Groups	4-45

5 Introduction to Oracle ACFS

Overview of Oracle ACFS	5-1
Understanding Oracle ACFS Concepts	5-3
About Oracle ACFS.....	5-4
About Oracle ACFS and Oracle Database Homes	5-4
About Oracle ASM Dynamic Volume Manager.....	5-5
About the Oracle ACFS Driver Model	5-6
About the Oracle ACFS Mount Model and Namespace	5-6
About the Oracle ACFS Mount Registry	5-6
About Oracle ACFS Snapshots.....	5-6
About Oracle ACFS and Backup and Restore.....	5-7
About Oracle ACFS Integration with Oracle ASM	5-7
Understanding Oracle ACFS Administration	5-8
Oracle ACFS and File Access and Administration Security	5-8
Oracle ACFS and Grid Infrastructure Installation	5-8
Oracle ACFS and Grid Infrastructure Configuration	5-8
Clusterware Resources and Oracle ACFS Administration	5-9
Oracle ACFS and Dismount or Shutdown Operations	5-10
Overview of Oracle ASM Dynamic Volume Manager	5-10

6 Using Views to Display Information

Views Containing Oracle ASM Disk Group Information	6-1
Viewing Oracle ASM File Access Control Information.....	6-4
Viewing Disk Region Information.....	6-5
Views Containing Oracle ACFS Information	6-6

7 Administering Oracle ASM Files, Directories, and Templates

What Types of Files Does Oracle ASM Support?	7-1
About Oracle ASM Filenames	7-2
Single File Creation Form	7-3
Multiple File Creation Form	7-4
Fully qualified File Name Form	7-4
Alias Oracle ASM Filename Forms	7-5
Creating a Tablespace in Oracle ASM: Using a Data File with an Alias Name	7-6
Alias Oracle ASM Filename with Template Form	7-6
Incomplete Oracle ASM Filename Form	7-7
Incomplete Oracle ASM Filename with Template Form	7-7
Creating and Referencing Oracle ASM Files in the Database	7-7
Creating Oracle ASM Files Using a Default File Location for Disk Group Specification	7-8
Using Oracle ASM Filenames in SQL Statements	7-9
Managing Alias Names for Oracle ASM Filenames	7-9
Adding an Alias Name for an Oracle ASM Filename	7-10
Renaming an Alias Name for an Oracle ASM Filename	7-10
Dropping an Alias Name for an Oracle ASM Filename	7-10
Dropping Files and Associated Aliases from a Disk Group	7-10
Managing Disk Group Directories	7-11
Creating a New Directory	7-11
Renaming a Directory	7-12
Dropping a Directory	7-12
Accessing Oracle ASM Files with the XML DB Virtual Folder	7-12
Inside /sys/asm	7-13
Using DBMS_FILE Transfer Utility for Oracle ASM	7-14
Managing Disk Group Templates	7-15
Template Attributes	7-15
Adding Templates to a Disk Group	7-17
Modifying a Disk Group Template	7-18
Dropping Templates from a Disk Group	7-18
Creating Tablespaces in Oracle ASM: Specifying Attributes with Templates	7-18

8 Performing Oracle ASM Data Migration With RMAN

Overview of Oracle ASM Data Migration	8-1
Purpose of Oracle ASM Data Migration	8-1
Basic Concepts of Oracle ASM Data Migration	8-2
Basics Steps of Data Migration to Oracle ASM Using RMAN	8-2
Preparing to Migrate the Database to Oracle ASM Using RMAN	8-3
Migrating the Database to Oracle ASM Using RMAN	8-5
Migrating a Database from Oracle ASM to Alternative Storage	8-9
Moving Data Files Between Oracle ASM Disk Groups Using RMAN	8-9

9 Administering Oracle ASM with Oracle Enterprise Manager

Oracle Automatic Storage Management Home Page	9-1
Accessing the Oracle ASM Home Page in Single-Instance Oracle Databases	9-2

Accessing the Oracle ASM Home Page in Oracle RAC Databases.....	9-3
Configuring Oracle ASM Initialization Parameters with Oracle Enterprise Manager	9-3
Bringing Disks Online and Offline	9-4
Managing Oracle ASM Users with Oracle Enterprise Manager	9-6
Managing Disk Groups with Oracle Enterprise Manager	9-7
Creating Disk Groups.....	9-7
Adding Disks to Disk Groups	9-9
Dropping Disks from Disk Groups	9-11
Dropping Disk Groups.....	9-12
Monitoring Disk Group Usage.....	9-13
Mounting and Dismounting Disk Groups	9-13
Administering Advanced Disk Group Properties	9-14
Configuring Disk Group Compatibility Attributes	9-15
Configuring Disk Repair Time.....	9-16
Configuring Smart Scan Compatibility	9-16
Configuring File Access Control.....	9-16
Managing Oracle ASM File Access Control with Oracle Enterprise Manager	9-16
Managing Directories, Files, and Aliases with Oracle Enterprise Manager.....	9-18
Managing Disk Group Templates with Oracle Enterprise Manager	9-20
Monitoring Oracle ASM Performance with Oracle Enterprise Manager.....	9-21
Monitoring General Oracle ASM Performance	9-21
Checking Disk Group Integrity.....	9-22
Backing Up Oracle ASM Files with Oracle Enterprise Manager.....	9-23
Performing Bad Block Recovery with Oracle Enterprise Manager	9-23
Migrating to Oracle ASM with Oracle Enterprise Manager.....	9-24
Oracle ASM Support Workbench	9-25
Oracle ASM: Oracle by Example Series.....	9-28
10 Administering Oracle ACFS With Oracle Enterprise Manager	
Managing Oracle ACFS with Oracle Enterprise Manager	10-1
Creating Oracle ACFS Volumes and File Systems	10-1
Viewing and Modifying Oracle ACFS Volumes and File Systems.....	10-3
Managing Oracle ACFS Snapshots with Oracle Enterprise Manager	10-5
11 Oracle ASM Configuration Assistant	
Oracle ASM Configuration Assistant GUI Tool	11-1
Starting Oracle ASM Configuration Assistant.....	11-1
Managing Oracle ASM Instances With Oracle ASM Configuration Assistant	11-2
Upgrading an Oracle ASM Instance with Oracle ASM Configuration Assistant.....	11-3
Managing Disk Groups With Oracle ASM Configuration Assistant	11-6
Managing Oracle ADVM Volumes With Oracle ASM Configuration Assistant.....	11-9
Managing Oracle ACFS File Systems With Oracle ASM Configuration Assistant	11-11
Creating an Oracle ACFS File System for a Database Home.....	11-13
Oracle ASM Configuration Assistant Command-Line Interface.....	11-14
Running ASMCA Command-Line	11-14
ASMCA Commands	11-15
Configure an Oracle ASM Instance	11-16

Upgrade an Oracle ASM Instance	11-16
Configure Parameters for an Oracle ASM Instance	11-17
Delete an Oracle ASM Instance.....	11-18
Create a Disk Group	11-18
Add a Disk to a Disk Group	11-20
Create a Volume.....	11-20
Create an Oracle ACFS File System	11-21

12 Oracle ASM Command-Line Utility

About ASMCMD.....	12-1
Types of ASMCMD Commands	12-1
About Oracle ASM Files, Filenames, Directories, and Aliases.....	12-2
System-Generated Filename or Fully Qualified Filename.....	12-2
Directory.....	12-2
Alias	12-3
Absolute Path and Relative Path	12-4
Wildcard Characters.....	12-5
Preparing to Run ASMCMD	12-6
Running ASMCMD in Interactive Mode.....	12-7
Running ASMCMD Without Any Options.....	12-7
Specifying the Type of Connection	12-8
Specifying the Verbose Mode.....	12-8
Including the Current Directory in the ASMCMD Prompt.....	12-8
Displaying the ASMCMD Version Number	12-8
Running ASMCMD in Noninteractive Mode	12-8
Getting Help.....	12-9
ASMCMD Instance Management Commands.....	12-9
dsget	12-10
dsset.....	12-10
lsct.....	12-11
lsop	12-12
lspwusr	12-12
orapwusr	12-13
shutdown.....	12-14
spbackup.....	12-14
spcopy	12-15
spget	12-16
spmmove.....	12-16
spset.....	12-17
startup	12-18
ASMCMD File Management Commands	12-18
cd	12-19
cp	12-19
du.....	12-21
find	12-21
ls.....	12-22
lsof	12-24

mkalias.....	12-25
mkdir.....	12-25
pwd.....	12-26
rm.....	12-26
rmalias.....	12-27
ASMCMD Disk Group Management Commands	12-28
chdg.....	12-28
chkdg.....	12-30
dropdg.....	12-31
iostat.....	12-31
lsattr.....	12-33
lsdg.....	12-34
lsdsk.....	12-36
lsod.....	12-38
md_backup.....	12-38
md_restore.....	12-39
mkdg.....	12-40
mount.....	12-42
offline.....	12-43
online.....	12-43
rebal.....	12-44
remap.....	12-45
setattr.....	12-46
umount.....	12-46
ASMCMD Template Management Commands	12-47
chtpl.....	12-47
lstmpl.....	12-48
mktpl.....	12-49
rmtmpl.....	12-49
ASMCMD File Access Control Commands	12-50
chgrp.....	12-50
chmod.....	12-51
chown.....	12-52
groups.....	12-53
grpmod.....	12-53
lsgrp.....	12-54
lsusr.....	12-55
mkgrp.....	12-55
mkusr.....	12-56
passwd.....	12-57
rmgrp.....	12-57
rmusr.....	12-58
ASMCMD Volume Management Commands	12-58
volcreate.....	12-59
voldelete.....	12-60
voldisable.....	12-61
volenable.....	12-62

volinfo	12-62
volresize.....	12-63
volset	12-64
volstat.....	12-65

13 Oracle ACFS Command-Line Tools

Basic Steps to Manage an Oracle ACFS	13-1
Creating an Oracle ACFS File System.....	13-1
Accessing an Oracle ACFS File System on a Different Node in the Cluster	13-4
Managing Oracle ACFS Snapshots.....	13-5
Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems.....	13-5
Deregistering an Oracle ACFS File System	13-5
Dismounting an Oracle ACFS File System.....	13-5
Disabling a Volume	13-6
Removing an Oracle ACFS File System and a Volume	13-6
Oracle ACFS Command-line Tools for Linux and UNIX Environments.....	13-7
fsck.....	13-7
mkfs.....	13-9
mount.....	13-10
umount	13-12
Oracle ACFS Command-line Tools for Windows Environments	13-13
acfschkdsk	13-13
acfsdismount.....	13-15
acfsformat	13-16
acfsmountvol	13-17
acfsutil detach.....	13-18
advmutil dismount	13-18
advmutil list.....	13-19
advmutil mount.....	13-20
Oracle ACFS Command-Line Tools for All Environments.....	13-22
acfsdbg	13-22
acfsutil info.....	13-25
acfsutil registry	13-27
acfsutil rmfs.....	13-28
acfsutil size	13-29
acfsutil snap create.....	13-30
acfsutil snap delete.....	13-32
acfsutil tune.....	13-32
advmutil tune	13-33
advmutil volinfo.....	13-34

A Creating Database Files in Oracle ASM Using SQL*Plus

Creating a Database that Uses Oracle ASM.....	A-1
Creating a Database that Uses Oracle ASM: Example.....	A-1
Creating Tablespaces in Oracle ASM.....	A-2
Creating Redo Logs in Oracle ASM.....	A-2

Creating Control Files in Oracle ASM	A-3
Creating Control Files in Oracle ASM: Example 1	A-3
Creating Control Files in Oracle ASM: Example 2	A-3
Creating Archive Log Files in Oracle ASM	A-4

B Oracle ACFS Advanced Topics

Limits of Oracle ACFS	B-1
Oracle ACFS Disk Space Usage.....	B-1
Oracle ACFS Error Handling	B-2
Oracle ACFS and NFS	B-2
Limits of Oracle ADVM	B-3
Oracle ACFS Drivers Resource Management	B-3
Oracle ACFS Registry Resource Management	B-4
Oracle ACFS Individual File System Resource Management	B-6
Oracle ACFS and Oracle Restart	B-6
Understanding Oracle ACFS I/O Failure Console Messages	B-7

Glossary

Index

List of Examples

4-1	Creating the DATA Disk Group.....	4-5
4-2	Creating the FRA Disk Group.....	4-6
4-3	Altering the DB_FILE_MULTIBLOCK_READ_COUNT Setting.....	4-8
4-4	Creating a Disk Group of 4K Sector Size.....	4-9
4-5	Using the QUORUM Keyword.....	4-11
4-6	Querying V\$ASM_DISK for Header Status.....	4-19
4-7	Setting Up Oracle ASM File Access Control.....	4-41
4-8	Using renamedg.....	4-47
6-1	Viewing Disk Group Attributes With V\$ASM_ATTRIBUTE.....	6-2
6-2	Viewing the Compatibility of a Disk Group with V\$ASM_DISKGROUP.....	6-3
6-3	Viewing Disks in Disk Groups with V\$ASM_DISK.....	6-3
6-4	Viewing Disk Group Clients With V\$ASM_CLIENT.....	6-3
6-5	Viewing Access Control Information in V\$ASM_USER.....	6-4
6-6	Viewing Access Control Information With V\$ASM_USERGROUP.....	6-4
6-7	Viewing Access Control Information With V\$ASM_USERGROUP_MEMBER.....	6-4
6-8	Viewing Access Control Information With V\$ASM_FILE.....	6-5
6-9	Viewing Intelligent Data Placement Information With V\$ASM_FILE.....	6-5
6-10	Viewing Intelligent Data Placement Information With V\$ASM_TEMPLATE.....	6-5
6-11	Viewing Snapshot Information in the V\$ASM_ACFSSNAPSHOTS View.....	6-6
6-12	Viewing Volume Information With V\$ASM_ACFSVOLUMES.....	6-6
6-13	Viewing Volume Information with V\$ASM_FILESYSTEM.....	6-7
6-14	Viewing Volume Information With V\$ASM_VOLUME.....	6-7
6-15	Viewing Volume Information With V\$ASM_VOLUME_STAT.....	6-8
7-1	Creating a Tablespace Using an Alias Name.....	7-6
7-2	Using an Oracle ASM File Name in a SQL Statement.....	7-9
7-3	Adding an Alias Name for an Oracle ASM Filename.....	7-10
7-4	Example: Renaming an Alias Name for an Oracle ASM Filename.....	7-10
7-5	Dropping an Alias Name for an Oracle ASM Filename.....	7-10
7-6	Dropping an Alias Name for an Oracle ASM Filename.....	7-10
7-7	Dropping Files and Associated Aliases from a Disk Group.....	7-11
7-8	Dropping Files and Associated Aliases from a Disk Group.....	7-11
7-9	Creating a New Directory.....	7-12
7-10	Creating a New Subdirectory.....	7-12
7-11	Renaming a Directory.....	7-12
7-12	Dropping a Directory.....	7-12
7-13	Adding a Template to a Disk Group.....	7-17
7-14	Adding a Template to a Disk Group.....	7-17
7-15	Modifying a Disk Group Template.....	7-18
7-16	Dropping a Template from a Disk Group.....	7-18
7-17	Using a User-Defined Template to Specify Attributes.....	7-18
8-1	Migrating the Online Redo Logs.....	8-8
11-1	Using asmca -silent -configureASM.....	11-16
11-2	Using asmca -silent -upgradeASM.....	11-17
11-3	Using asmca -silent -configureParameter.....	11-17
11-4	Using asmca -silent -deleteASM.....	11-18
11-5	Using asmca -silent -createDiskGroup.....	11-19
11-6	Using asmca -silent -addDisk.....	11-20
11-7	Using asmca -silent -createVolume.....	11-21
11-8	Using asmca -silent -createACFS.....	11-22
12-1	Using Wildcards With ASMCMD.....	12-5
12-2	Running ASMCMD in Noninteractive Mode.....	12-8
12-3	Displaying ASMCMD Help Text.....	12-9
12-4	Using dsget.....	12-10
12-5	Using dsset.....	12-11

12-6	Using lsct.....	12-12
12-7	Using lsop.....	12-12
12-8	Using lspwusr.....	12-13
12-9	Using orapwusr.....	12-13
12-10	Using shutdown.....	12-14
12-11	Using spbackup.....	12-15
12-12	Using spcopy.....	12-16
12-13	Using spget.....	12-16
12-14	Using spmove.....	12-17
12-15	Using spset.....	12-18
12-16	Using startup.....	12-18
12-17	Using cd.....	12-19
12-18	Using cp.....	12-20
12-19	Using du.....	12-21
12-20	Using find.....	12-22
12-21	Using ls.....	12-23
12-22	Using lsof.....	12-24
12-23	Using mkalias.....	12-25
12-24	Using mkdir.....	12-26
12-25	Using pwd.....	12-26
12-26	Using rm.....	12-27
12-27	Using rmalias.....	12-28
12-28	Tags for the chdg XML Configuration Template.....	12-29
12-29	chdg sample XML configuration file.....	12-30
12-30	Using chdg.....	12-30
12-31	Using chkdg.....	12-31
12-32	Using dropdg.....	12-31
12-33	Using iostat.....	12-32
12-34	Using lsattr.....	12-34
12-35	Using lsdg.....	12-35
12-36	Using lsdsk.....	12-37
12-37	Using lsod.....	12-38
12-38	Using md_backup.....	12-39
12-39	Using md_restore.....	12-40
12-40	Tags for mkdg XML Configuration File.....	12-41
12-41	mkdg sample XML configuration file.....	12-42
12-42	Using mkdg.....	12-42
12-43	Using mount.....	12-43
12-44	Using offline.....	12-43
12-45	Using online.....	12-44
12-46	Using rebal.....	12-45
12-47	Using remap.....	12-45
12-48	Using setattr.....	12-46
12-49	Using umount.....	12-47
12-50	Using chtmpl.....	12-48
12-51	Using lstmpl.....	12-48
12-52	Using mktmpl.....	12-49
12-53	Using rmtmpl.....	12-50
12-54	Using chgrp.....	12-51
12-55	Using chmod.....	12-52
12-56	Using chown.....	12-53
12-57	Using groups.....	12-53
12-58	Using grpmod.....	12-54
12-59	Using lsgrp.....	12-55
12-60	Using lsusr.....	12-55

12-61	Using mkgrp	12-56
12-62	Using mkusr	12-56
12-63	Using passwd	12-57
12-64	Using rmgrp.....	12-58
12-65	Using rmusr	12-58
12-66	Using volcreate.....	12-60
12-67	Using voldelete.....	12-61
12-68	Using voldisable.....	12-62
12-69	Using volenable.....	12-62
12-70	Using volinfo	12-63
12-71	Using volresize	12-64
12-72	Using volset	12-64
12-73	Using volstat	12-65
13-1	Using fsck	13-9
13-2	Using mkfs	13-10
13-3	Using mount	13-12
13-4	Using umount.....	13-12
13-5	Using acfschkdsk	13-15
13-6	Using acfsdismount	13-15
13-7	Using acfsformat	13-16
13-8	Using acfsmountvol.....	13-17
13-9	Using acfsutil detach	13-18
13-10	Using advmutil dismount.....	13-19
13-11	Using advmutil list	13-20
13-12	Using advmutil mount.....	13-21
13-13	Using acfsdbg	13-24
13-14	Using acfsutil info	13-26
13-15	Using acfsutil registry	13-28
13-16	Using acfsutil rmfs	13-29
13-17	Using acfsutil size	13-30
13-18	Using acfsutil snap create	13-31
13-19	Using acfsutil snap delete	13-32
13-20	Using acfsutil tune	13-33
13-21	Using advmutil tune.....	13-34
13-22	Using advmutil volinfo	13-35
B-1	Loading Drivers	B-3
B-2	Mounting File Systems.....	B-4

List of Tables

3-1	One Operating System Group and One Set of Privileges for All Oracle ASM Users ...	3-24
3-2	Separated Operating System Groups and Privileges for Oracle ASM Users.....	3-25
4-1	Mirroring Options for Oracle ASM Disk Group Types	4-23
4-2	Examples of Disk Group Compatibility Attribute Settings	4-34
4-3	Features Enabled by Disk Group Compatibility Attribute Settings.....	4-36
4-4	Maximum Oracle ASM File Sizes for Disk Groups With AU_SIZE Equal to 1 MB.....	4-38
6-1	Oracle ASM Dynamic Views for Disk Group Information.....	6-1
6-2	Dynamic View for Oracle ACFS Information.....	6-6
7-1	File Types Supported by Automatic Storage Management.....	7-1
7-2	Valid Usage for Oracle ASM Filename Forms.....	7-3
7-3	Oracle File Types and Oracle ASM File Type Tags.....	7-4
7-4	Permitted Values for Oracle ASM Template Striping Attribute.....	7-16
7-5	Permitted Values for Oracle ASM Template Redundancy Attribute.....	7-16
7-6	Oracle ASM System Default Templates Attribute Settings	7-16
7-7	Permitted Values for Oracle ASM Template Primary and Mirror Region Attribute...	7-17
11-1	Options for the asmca command-line tool	11-15
11-2	Options for the -configureASM command.....	11-16
11-3	Options for the -configureParameter command	11-17
11-4	Options for the -deleteASM command.....	11-18
11-5	Options for the -createDiskGroup command	11-19
11-6	Options for the -addDisk command	11-20
11-7	Options for the -createVolume command	11-21
11-8	Options for the -createACFS command.....	11-22
12-1	Types of ASMCMD Commands	12-2
12-2	ASMCMD Interactive Options.....	12-7
12-3	Summary of ASMCMD Instance Management Commands	12-9
12-4	Options for the dsget command	12-10
12-5	Options for the dsset command.....	12-11
12-6	Options for the lsct command.....	12-11
12-7	Options for the lspwusr Command	12-13
12-8	Options for the orapwusr Command.....	12-13
12-9	Options for the shutdown command.....	12-14
12-10	Options for the spbackup command.....	12-14
12-11	Options for the spcopy command	12-15
12-12	Options for the spmove command.....	12-17
12-13	Options for the spset command.....	12-17
12-14	Options for the startup command	12-18
12-15	Summary of ASMCMD File Management Commands.....	12-18
12-16	Options for the cd Command.....	12-19
12-17	Options for the cp Command.....	12-20
12-18	Options for the du command	12-21
12-19	Options for the find Command	12-21
12-20	Options for the ls command	12-22
12-21	Options for the lsdf command	12-24
12-22	Options for the mkalias command	12-25
12-23	Options for the mkdir command.....	12-25
12-24	Options for the rm command.....	12-26
12-25	Options for the rmalias command.....	12-27
12-26	Summary of ASMCMD Disk Group Management Commands	12-28
12-27	Options for the chdg Command	12-29
12-28	Options for the chkdg Command.....	12-30
12-29	Options for the dropdg Command	12-31
12-30	Options for the iostat command	12-32

12-31	Attribute descriptions for iostat command output.....	12-32
12-32	Options for the lsattr command.....	12-33
12-33	Options for the lsdg command.....	12-34
12-34	Attribute descriptions for lsdg command output.....	12-35
12-35	Options for the lsdsd command.....	12-36
12-36	Options for the lsod command.....	12-38
12-37	Options for the md_backup Command.....	12-39
12-38	Options for the md_restore command.....	12-40
12-39	Options for the mkgd Command.....	12-41
12-40	Options for the mount Command.....	12-42
12-41	Options for the offline command.....	12-43
12-42	Options for the online command.....	12-44
12-43	Options for the rebal command.....	12-44
12-44	Options for the remap command.....	12-45
12-45	Options for the setattr command.....	12-46
12-46	Options for the umount Command.....	12-46
12-47	Summary of ASMCMD Template Management Commands.....	12-47
12-48	Options for the chtmpl Command.....	12-47
12-49	Options for the lstmpl command.....	12-48
12-50	Options for the mktmpl Command.....	12-49
12-51	Options for the rmtmpl Command.....	12-49
12-52	Summary of ASMCMD File Access Control Commands.....	12-50
12-53	Options for the chgrp Command.....	12-50
12-54	Options for the chmod Command.....	12-51
12-55	Options for the chown Command.....	12-52
12-56	Options for the groups Command.....	12-53
12-57	Options for the grpmod Command.....	12-54
12-58	Options for the lsgrp command.....	12-54
12-59	Options for the lsusr command.....	12-55
12-60	Options for the mkgrp Command.....	12-56
12-61	Options for the mkusr command.....	12-56
12-62	Options for the passwd command.....	12-57
12-63	Options for the rmgrp Command.....	12-57
12-64	Options for the rmusr command.....	12-58
12-65	Summary of ASMCMD Volume Management Commands.....	12-59
12-66	Options for the volcreate command.....	12-59
12-67	Options for the voldelete command.....	12-61
12-68	Options for the voldisable command.....	12-61
12-69	Options for the volenable command.....	12-62
12-70	Options for the volinfo command.....	12-63
12-71	Options for the volresize command.....	12-63
12-72	Options for the volset command.....	12-64
12-73	Options for the volstat command.....	12-65
13-1	Summary of ACFS Commands for Linux and UNIX.....	13-7
13-2	Options for the fsck command.....	13-8
13-3	Options for the mkfs command.....	13-9
13-4	Options for the mount command.....	13-10
13-5	Options for the umount command.....	13-12
13-6	Summary of ACFS File System Commands for Windows.....	13-13
13-7	Summary of Non-ACFS File System Commands for Windows.....	13-13
13-8	Options for the acfschkdsk command.....	13-13
13-9	Options for the acfsdismount command.....	13-15
13-10	Options for the acfsformat command.....	13-16
13-11	Options for the acfsmountvol command.....	13-17
13-12	Options for the advmutil dismount command.....	13-18

13-13	Options for the advmutil list command	13-20
13-14	Options for the advmutil mount command	13-20
13-15	Summary of ACFS Commands for All Environments	13-22
13-16	Options for the acfsdbg command	13-22
13-17	Subcommands for acfsdbg	13-23
13-18	Options for the acfsutil info command	13-25
13-19	Options for the acfsutil registry command	13-27
13-20	Options for the acfsutil rmfs command	13-29
13-21	Options for the acfsutil size command	13-29
13-22	Options for the acfsutil snap create command	13-31
13-23	Options for the acfsutil snap delete command	13-32
13-24	Options for the acfsutil tune command	13-33
13-25	Options for the advmutil tune command	13-33
13-26	Options for the advmutil volinfo command	13-34

Preface

The *Oracle Database Storage Administrator's Guide* describes how to administer storage for Oracle databases. This Preface contains the following topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Related Documents](#)
- [Conventions](#)

Audience

The audience for this book includes system administrators, database administrators, and storage administrators. The *Oracle Database Storage Administrator's Guide* is intended for database and storage administrators who perform the following tasks:

- Administer and manage storage
- Configure and administer ASM

To use this document, you should be familiar with basic Oracle Database concepts and administrative procedures. Also, you might want to review the documentation for Oracle Clusterware and Oracle Real Application Clusters (Oracle RAC). See "[Related Documents](#)" on page xxii.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at

<http://www.fcc.gov/cgb/consumerfacts/trs.html>, and a list of phone numbers is available at <http://www.fcc.gov/cgb/dro/trsphonebk.html>.

Related Documents

For more information, refer to the following Oracle resources:

- *Oracle Database 2 Day DBA*
- *Oracle Database 2 Day + Real Application Clusters Guide*
- *Oracle Database Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle Database Net Services Administrator's Guide*
- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Real Application Clusters Administration and Deployment Guide*
- Platform-specific Oracle Database, Oracle Clusterware, and Oracle RAC installation guides

To download release notes, installation documentation, white papers, or other collateral, refer to the Oracle Technology Network (OTN). You must register online before using OTN; registration is free and can be done at

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

If you have a username and password for OTN, then you can go directly to the documentation section of the OTN Web site at

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

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

What's New in Oracle Database Storage?

This section describes the new storage features for Oracle Database 11g and it contains the following topic:

- [New Storage Features for Oracle Database 11g Release 2 \(11.2\)](#)
- [New Storage Features for Oracle Database 11g Release 1 \(11.1\)](#)

See Also: *Oracle Database New Features Guide* for a complete description of the new features in Oracle Database 11g release 2 (11.2)

New Storage Features for Oracle Database 11g Release 2 (11.2)

This section describes the following Oracle Database 11g release 2 (11.2) Oracle Automatic Storage Management (Oracle ASM) features:

- Oracle Automatic Storage Management Cluster File System (Oracle ACFS)

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a new multi-platform, scalable file system, and storage management design that extends Oracle ASM technology to support of all your application data in both single host and cluster configurations. Additionally, Oracle ACFS provides snapshot functionality for a point in time copy of an Oracle ACFS system.

For more information, see "[Overview of Oracle ACFS](#)" on page 5-1.

- Oracle ASM Dynamic Volume Manager

The Oracle ASM Dynamic Volume Manager (Oracle ADVM) provides volume management services and a standard disk device driver interface to clients. File systems and other disk-based applications issue I/O requests to Oracle ADVM volume devices as they would to other storage devices on a vendor operating system.

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is layered on Oracle ASM through the Oracle ADVM interface. With the addition of the Oracle ADVM, Oracle ASM becomes a complete storage solution of your data for both database and non-database file needs.

For more information, see "[Overview of Oracle ASM Dynamic Volume Manager](#)" on page 5-10.

- Oracle Automatic Storage Management Cluster File System Snapshots

An Oracle ACFS Snapshot is a point-in-time copy of a file system that can provide up to 63 snapshot images. Oracle ACFS Snapshot performs fast creation of

persistent Oracle ACFS images at a specific point in time with very low overhead leveraging the Copy-on-Write (COW) technology.

Even as the file system changes, the snapshot does not, giving you the ability to view the file system as it was at the time the snapshot was created. Initially, snapshots are read-only, which preserves the point-in-time they captured.

For more information, see ["About Oracle ACFS Snapshots"](#) on page 5-6.

- ASMCMD Enhancements

This feature extends ASMCMD to provide full functionality, so that any operation that can be performed with SQL commands can be performed with ASMCMD. The added functionality includes the following areas:

- Disk, failure group, and disk group operations
- Disk group attribute operations
- Oracle ASM File Access Control user and group operations
- Template operations
- Oracle ASM instance operations
- File operations
- Oracle ASM volume operations

Changes were also made to standardize ASMCMD command-line and command options.

For more information about ASMCMD commands, see ["Types of ASMCMD Commands"](#) on page 12-1.

- Intelligent Data Placement

The Intelligent Data Placement feature enables you to specify disk regions on Oracle ASM disks to ensure that frequently accessed data is placed on the outermost (hot) tracks which provide higher performance.

For more information, see ["Intelligent Data Placement"](#) on page 4-15.

- Specifying the Sector Size for Disk Drives

Oracle ASM provides the capability to specify a sector size of 512 bytes or 4096 kilobytes with the `SECTOR_SIZE` disk group attribute when creating disk groups.

Oracle ASM provides support for 4 KB sector disk drives without a performance penalty.

For information about setting the sector size for disk drives, see ["Specifying the Sector Size for Drives"](#) on page 4-8.

See Also: *Oracle Database SQL Language Reference* for information about the disk group attributes and the `CREATE DISKGROUP` SQL statement

- Disk Group Rename

The `renamedg` tool enables you to change the name of a cloned disk group.

For information about renaming disk groups with `renamedg`, see ["Renaming Disks Groups"](#) on page 4-45.

- Oracle ASM File Access Control

Oracle ASM on Linux and UNIX platforms implements access control on its files to isolate different database instances from each other and prevent accidental access that could lead to the corruption of data files. Oracle ASM implements new SQL statements and ASMCMD commands to grant, modify, and deny file permissions. The new security model and syntax is consistent with those that have been implemented for the objects represented in an Oracle database.

This feature allows multiple database instances to store their Oracle ASM files in the same disk group and be able to consolidate multiple databases while preventing an unauthorized database instance from accessing and overwriting files of a different instance.

For more information about Oracle ASM File Access Control, see "[Managing Oracle ASM File Access Control for Disk Groups](#)" on page 4-38, "[Managing Oracle ASM File Access Control with Oracle Enterprise Manager](#)" on page 9-16, and "[ASMCMD File Access Control Commands](#)" on page 12-50.

- Oracle Cluster Registry (OCR) and Voting Files stored in Oracle ASM

This feature enables storage of Oracle Cluster Registry (OCR) and the voting file in Oracle ASM disk groups. This enables Oracle ASM to provide a unified storage solution, storing all the data for the clusterware and the database, without the need for third-party volume managers or cluster file systems.

For more information, see "[Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#)" on page 4-10.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about administering voting files and Oracle Cluster Registry (OCR)

- Oracle Restart

Oracle Restart enables the startup and automatic restart of single-instance Oracle Database through the Cluster Ready Services (CRS) component of Oracle Clusterware on a single server. This provides higher availability and automated management of single-instance Oracle Database and Oracle ASM instances.

For more information, see "[Using Oracle Restart](#)" on page 3-12.

See Also:

- *Oracle Database Administrator's Guide* for information about configuring and administering Oracle Restart
- *Oracle Real Application Clusters Administration and Deployment Guide* for information about automatically restarting single-instance databases residing on an Oracle RAC node

- Oracle ASM Configuration Assistant (ASMCA)

Oracle ASM Configuration Assistant provides a GUI interface for installing and configuring Oracle ASM instances, disk groups, volumes, and file systems.

In addition, the ASMCA command-line interface provides functionality for installing and configuring Oracle ASM instances, disk groups, volumes, and file systems in a non-GUI utility.

For more information about using Oracle ASM Configuration Assistant, see [Chapter 11, "Oracle ASM Configuration Assistant"](#).

- Oracle Enterprise Manager Support of Oracle ACFS

Oracle Enterprise Manager has enhanced Oracle ASM support to enable you to manage Oracle Automatic Storage Management Cluster File System (Oracle ACFS) technology and Oracle ASM Dynamic Volume Manager. These features operate as a comprehensive management solution to extend Oracle ASM technology to support all your application data files in database and non-database and in both single host and cluster configurations.

Oracle Enterprise Manager supports for the following:

- Create, modify and delete Oracle ACFS
- Monitor space usage and performance features of Oracle ACFS
- Support Oracle ASM Dynamic Volume Manager features as required by Oracle ACFS

For more information about using Oracle Enterprise Manager to administer Oracle ASM, see [Chapter 10, "Administering Oracle ACFS With Oracle Enterprise Manager"](#).

- Oracle Enterprise Manager Database Control Enhancements for Improved Oracle ASM Manageability

Oracle Enterprise Manager has been enhanced to manage Oracle ASM features, such as Oracle ASM File Access Control, OCR and voting disks in Oracle ASM, disk resync, Oracle ASM rolling migrations, disk regions, and Oracle ASM manageability and infrastructure.

For more information about using Oracle Enterprise Manager to administer Oracle ASM, see [Chapter 9, "Administering Oracle ASM with Oracle Enterprise Manager"](#).

- Oracle Enterprise Manager Oracle ASM Support Workbench

Oracle Enterprise Manager Oracle ASM Support Workbench has been enhanced to help diagnose and package incidents to Oracle Support Services for Oracle ASM instances. This feature extends the benefit of Oracle Enterprise Manager Support Workbench to Oracle ASM by helping you package all necessary diagnostic data for incidents.

For more information about using Oracle Enterprise Manager to administer Oracle ASM, see ["Oracle ASM Support Workbench"](#) on page 9-25.

See Also:

- *Oracle Database Administrator's Guide* for information about viewing problems with the Oracle Enterprise Manager Support Workbench
- *Oracle Database 2 Day DBA* for information about accessing the Oracle Enterprise Manager Support Workbench home page and viewing problems using the Oracle Enterprise Manager Support Workbench

New Storage Features for Oracle Database 11g Release 1 (11.1)

This section describes the following Oracle Database 11g release 1 (11.1) Oracle Automatic Storage Management (Oracle ASM) features:

- *Oracle Database Storage Administrator's Guide* is New for this release

This book, the *Oracle Database Storage Administrator's Guide*, is new for Oracle Database 11g release 1 (11.1) and it is the primary information source for Oracle Automatic Storage Management features.

- Oracle ASM Fast Mirror Resync

Oracle ASM fast mirror resync quickly resynchronizes Oracle ASM disks within a disk group after transient disk path failures provided that the disk drive media is not corrupted. Any failures that render a failure group temporarily unavailable are considered transient failures. Disk path malfunctions, such as cable disconnections, host bus adapter or controller failures, or disk power supply interruptions, can cause transient failures.

The duration of a fast mirror resync depends on the duration of the outage. The duration of a resynchronization is typically much shorter than the amount of time required to completely rebuild an entire Oracle ASM disk group.

For more information about Oracle ASM fast mirror resync, see "[Oracle ASM Fast Mirror Resync](#)" on page 4-27.

- Oracle ASM Rolling Upgrade

You can now place an Oracle ASM Cluster in rolling upgrade mode, which enables you to operate with mixed Oracle ASM versions starting with Oracle Database 11g release 1 (11.1) and later. As a result, Oracle ASM nodes can be independently upgraded or patched without affecting database availability.

For more information about Oracle ASM rolling upgrade, see "[Using Oracle ASM Rolling Upgrade](#)" on page 3-21.

- New SYSASM Privilege and OSASM operating system group for Oracle ASM Administration

This feature introduces a new SYSASM privilege that is specifically intended for performing Oracle ASM administration tasks. Using the SYSASM privilege instead of the SYSDBA privilege provides a clearer division of responsibility between Oracle ASM administration and database administration.

OSASM is a new operating system group that is used exclusively for Oracle ASM. Members of the OSASM group can connect as SYSASM using operating system authentication and have full access to Oracle ASM.

For more information about the SYSASM privilege, see "[Authentication for Accessing Oracle ASM Instances](#)" on page 3-22.

- Oracle ASM Scalability and Performance Enhancements

Oracle ASM file extent management has been enhanced to improve performance and to use significantly less SGA memory to store file extents. When Oracle ASM files increase in size, the size of each new extent also increases automatically so that fewer extent pointers are required to describe the file. This feature improves performance when accessing Oracle ASM files that are 20 GB and larger, up to 128 TB. Very large databases (VLDBs) often require these large file sizes.

In addition, when you create new disk groups you now have multiple allocation unit (AU) size options, such as 1, 2, 4, 8, 16, 32, and 64. You might obtain significant performance improvements by selecting larger allocation units depending on the type of workloads (typically large sequential I/O) and storage system types.

For more information about scalability and performance enhancements, see "[Extents](#)" on page 1-8.

- Oracle ASM Command Line Utility (ASMCMD) Commands and Options
ASMCMD has the following four new commands: `lsdisk`, `md_backup`, `md_restore` and `remap`. In addition, you can use new options for the `ls` and `lsdg` commands. The following describes the four new Oracle ASM commands:

For more information about new and enhanced ASMCMD commands, "[Types of ASMCMD Commands](#)" on page 12-1.

- Enhancements to SQL statements to manage Oracle ASM disk groups
The `ALTER DISKGROUP`, `CREATE DISKGROUP`, and `DROP DISKGROUP` SQL statements have been enhanced with additional syntax options, including:
 - Syntax that sets various attributes of a disk group
 - Syntax for checking the consistency of disk groups, disks, and files in an Oracle ASM environment
 - Syntax options when mounting a disk group.
 - Syntax that takes a disk offline for repair and then bring it back online.
 - Syntax that drops a disk group that can no longer be mounted by an Oracle ASM instance.

For more information about administering disk groups with the `ALTER DISKGROUP`, `CREATE DISKGROUP`, and `DROP DISKGROUP` SQL statements, see [Chapter 4, "Administering Oracle ASM Disk Groups"](#).

See Also: *Oracle Database SQL Language Reference* for information about the disk group attributes and the `CREATE DISKGROUP` SQL statement

- New Attributes for Disk Group Compatibility
To enable some of the new Oracle ASM features, you can use two new disk group compatibility attributes, `COMPATIBLE.RDBMS` and `COMPATIBLE.ASM`. These attributes specify the minimum software version that is required to use disk groups for the database and for Oracle ASM respectively. This feature enables heterogeneous environments with disk groups from both Oracle Database 10g and Oracle Database 11g. By default, both attributes are set to 10.1. You must advance these attributes to take advantage of the new features.

For more information about disk group compatibility, see "[Disk Group Compatibility](#)" on page 4-31.

- Oracle ASM Preferred Read Failure Groups
This feature is useful in extended clusters where remote nodes have asymmetric access with respect to performance. This enables more efficient use of network resources by eliminating the need to use the network to perform read operations.
Oracle ASM in Oracle Database 10g always reads the primary copy of a mirrored extent set. In Oracle Database 11g, when you configure Oracle ASM failure groups it might be more efficient for a node to read from a failure group that is closest to the node, even if it is a secondary extent. You can configure your database to read from a particular failure group extent by configuring preferred read failure groups.

For more information about preferred read failure groups, see "[Preferred Read Failure Groups](#)" on page 4-29.

- Oracle ASM Fast Rebalance

Rebalance operations that occur while a disk group is in RESTRICTED mode eliminate the lock and unlock extent map messaging between Oracle ASM instances in Oracle RAC environments, improving overall rebalance throughput.

For more information about Oracle ASM fast rebalance, see ["About Restricted Mode"](#) on page 3-15.

Introduction to Oracle Automatic Storage Management

This chapter describes Oracle Automatic Storage Management (Oracle ASM) concepts and provides an overview of Oracle ASM features. This chapter contains the following topics:

- [Overview of Oracle Automatic Storage Management](#)
- [Understanding Oracle ASM Concepts](#)
- [Understanding Oracle ASM Disk Group Administration](#)

For a list of the terms that are used in the *Oracle Database Storage Administrator's Guide* and their definitions, refer to the [Glossary](#) in this guide.

See Also: <http://www.oracle.com/technology/asm/> for more information about Oracle ASM

Overview of Oracle Automatic Storage Management

Oracle ASM is a **volume manager** and a **file system** for Oracle database files that supports single-instance Oracle Database and Oracle Real Application Clusters (Oracle RAC) configurations. Oracle ASM is Oracle's recommended storage management solution that provides an alternative to conventional volume managers, file systems, and raw devices.

Oracle ASM uses **disk groups** to store data files; an Oracle ASM disk group is a collection of disks that Oracle ASM manages as a unit. Within a disk group, Oracle ASM exposes a file system interface for Oracle database files. The content of files that are stored in a disk group is evenly distributed to eliminate hot spots and to provide uniform performance across the disks. The performance is comparable to the performance of raw devices.

You can add or remove disks from a disk group while a database continues to access files from the disk group. When you add or remove disks from a disk group, Oracle ASM automatically redistributes the file contents and eliminates the need for downtime when redistributing the content. For information about administering disk groups, see [Chapter 4, "Administering Oracle ASM Disk Groups"](#).

The Oracle ASM volume manager functionality provides flexible server-based mirroring options. The Oracle ASM normal and high redundancy disk groups enable two-way and three-way mirroring respectively. You can use external redundancy to enable a Redundant Array of Independent Disks (RAID) storage subsystem to perform the mirroring protection function.

Oracle ASM also uses the Oracle Managed Files (OMF) feature to simplify database file management. OMF automatically creates files in designated locations. OMF also names files and removes them while relinquishing space when tablespaces or files are deleted.

Oracle ASM reduces the administrative overhead for managing database storage by consolidating data storage into a small number of disk groups. The smaller number of disk groups consolidates the storage for multiple databases and provides for improved I/O performance.

Oracle ASM files can coexist with other storage management options such as raw disks and third-party file systems. This capability simplifies the integration of Oracle ASM into pre-existing environments.

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a multi-platform, scalable file system, and storage management technology that extends Oracle ASM functionality to support customer files maintained outside of the Oracle Database. The Oracle ASM Dynamic Volume Manager (Oracle ADVM) provides volume management services and a standard disk device driver interface to clients. For information about Oracle ACFS and Oracle ADVM, see [Chapter 5, "Introduction to Oracle ACFS"](#).

Oracle Enterprise Manager includes a wizard that enables you to migrate non-Oracle ASM database files to Oracle ASM. Oracle ASM also has easy to use management interfaces such as SQL*Plus, the ASMCMD command-line interface, Oracle ASM Configuration Assistant, and Oracle Enterprise Manager. For information about using Oracle Enterprise Manager, see [Chapter 9, "Administering Oracle ASM with Oracle Enterprise Manager"](#). For information about the ASMCMD command-line interface, see [Chapter 12, "Oracle ASM Command-Line Utility"](#).

See Also: *Oracle Database Administrator's Guide* for information about the Oracle database structure and storage

Understanding Oracle ASM Concepts

This section describes concepts for the key Oracle ASM components and it contains the following topics:

- [About Oracle ASM Instances](#)
- [About Oracle ASM Disk Groups](#)
- [About Mirroring and Failure Groups](#)
- [About Oracle ASM Disks](#)
- [About Oracle ASM Files](#)

For information about preparing your storage environment, see [Chapter 2, "Considerations for Oracle ASM Storage"](#).

About Oracle ASM Instances

An Oracle ASM instance is built on the same technology as an Oracle Database instance. An Oracle ASM instance has a System Global Area (SGA) and background processes that are similar to those of Oracle Database. However, because Oracle ASM performs fewer tasks than a database, an Oracle ASM SGA is much smaller than a database SGA. In addition, Oracle ASM has a minimal performance effect on a server. Oracle ASM instances mount disk groups to make Oracle ASM files available to database instances; Oracle ASM instances do not mount databases.

Oracle ASM is installed in the Oracle grid infrastructure home before the Oracle Database is installed in a separate Oracle home.

Oracle ASM and database instances require shared access to the disks in a disk group. Oracle ASM instances manage the metadata of the disk group and provide file layout information to the database instances.

Oracle ASM metadata is the information that Oracle ASM uses to control a disk group and the metadata resides within the disk group. Oracle ASM metadata includes the following information:

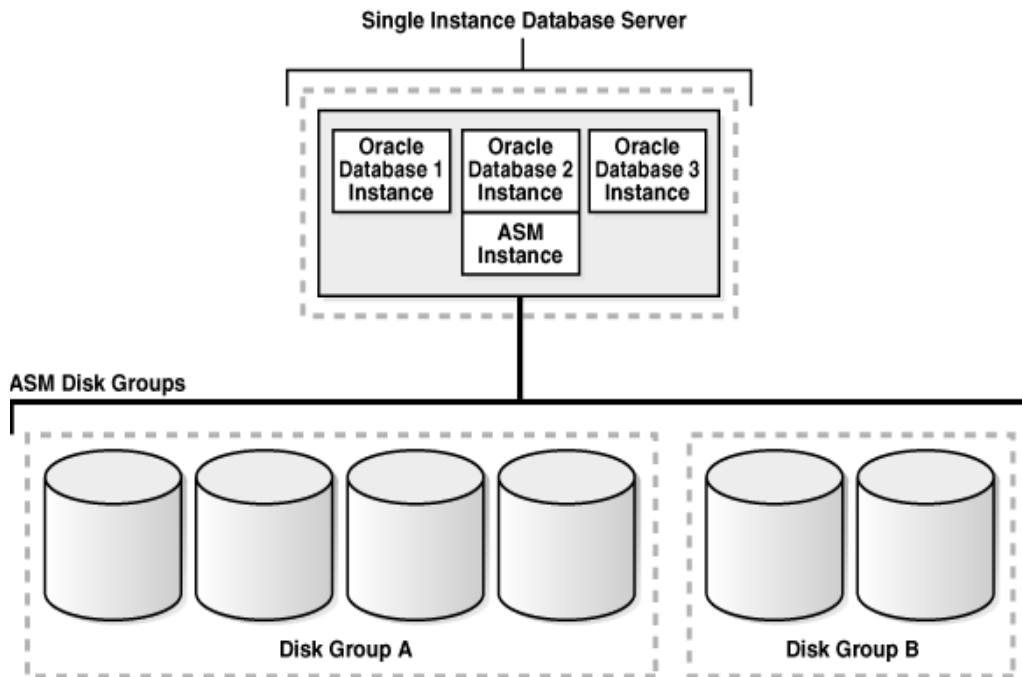
- The disks that belong to a disk group
- The amount of space that is available in a disk group
- The filenames of the files in a disk group
- The location of disk group data file **extents**
- A redo log that records information about atomically changing metadata blocks
- Oracle ADVM volume information

Oracle ASM instances can be clustered using Oracle Clusterware; there is one Oracle ASM instance for each cluster node. If there are several database instances for different databases on the same node, then the database instances share the same single Oracle ASM instance on that node.

If the Oracle ASM instance on a node fails, then all of the database instances on that node also fail. Unlike a file system driver failure, an Oracle ASM instance failure does not require restarting the operating system. In an Oracle RAC environment, the Oracle ASM and database instances on the surviving nodes automatically recover from an Oracle ASM instance failure on a node.

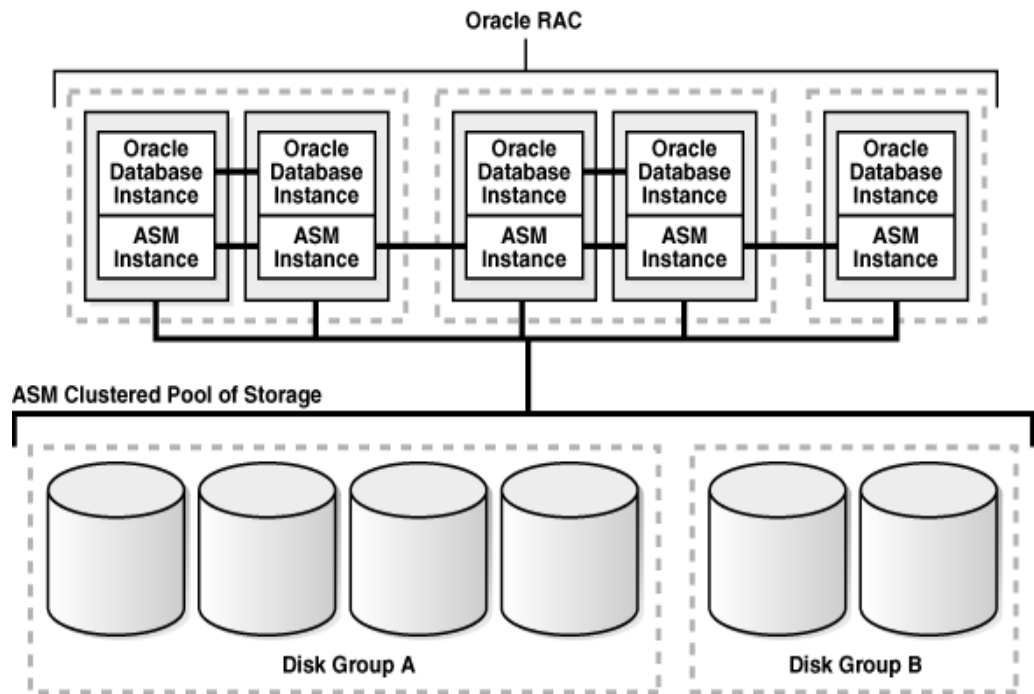
[Figure 1–1](#) shows a single node configuration with one Oracle ASM instance and multiple database instances. The Oracle ASM instance manages the metadata and provides space allocation for the Oracle ASM files. When a database instance creates or opens an Oracle ASM file, it communicates those requests to the Oracle ASM instance. In response, the Oracle ASM instance provides file extent map information to the database instance.

Figure 1–1 Oracle ASM for Single-Instance Oracle Databases



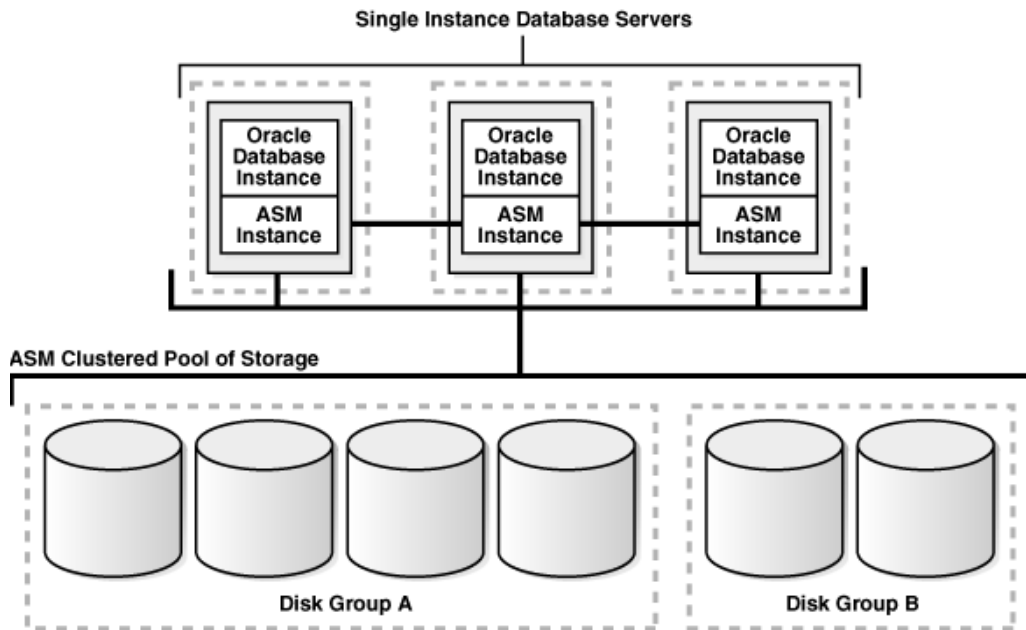
In [Figure 1–1](#), there are two disk groups: one disk group has four disks and the other has two disks. The database can access both disk groups. The configuration in [Figure 1–1](#) shows multiple database instances, but only one Oracle ASM instance is needed to serve the multiple database instances.

[Figure 1–2](#) shows an Oracle ASM cluster in an Oracle RAC environment where Oracle ASM provides a clustered pool of storage. There is one Oracle ASM instance for each node serving multiple Oracle RAC or single-instance databases in the cluster. All of the databases are consolidated and sharing the same two Oracle ASM disk groups.

Figure 1–2 Oracle ASM Cluster Configuration with Oracle RAC

A clustered storage pool can be shared by multiple single-instance Oracle Databases as shown in [Figure 1–3](#). In this case, multiple databases share common disk groups. A shared Oracle ASM storage pool is achieved by using Oracle Clusterware. However, in such environments an Oracle RAC license is not required.

To share a disk group among multiple nodes, you must install Oracle Clusterware on all of the nodes, regardless of whether you install Oracle RAC on the nodes. Oracle ASM instances that are on separate nodes do not need to be part of an Oracle ASM cluster. However, if the Oracle ASM instances are not part of an Oracle ASM cluster, they cannot communicate with each other. Multiple nodes that are not part of an Oracle ASM cluster cannot share a disk group.

Figure 1-3 Oracle ASM Cluster with Single-Instance Oracle Databases

About Oracle ASM Disk Groups

A disk group consists of multiple disks and is the fundamental object that Oracle ASM manages. Each disk group contains the metadata that is required for the management of space in the disk group. Disk group components include disks, files, and allocation units.

Files are allocated from disk groups. Any Oracle ASM file is completely contained within a single disk group. However, a disk group might contain files belonging to several databases and a single database can use files from multiple disk groups. For most installations you need only a small number of disk groups, usually two, and rarely more than three.

About Mirroring and Failure Groups

Mirroring protects data integrity by storing copies of data on multiple disks. When you create a disk group, you specify an Oracle ASM disk group type based on one of the following three redundancy levels:

- **Normal** for 2-way mirroring
- **High** for 3-way mirroring
- **External** to not use Oracle ASM mirroring, such as when you configure hardware RAID for redundancy

The redundancy level controls how many disk failures are tolerated without dismounting the disk group or losing data. The disk group type determines the mirroring levels with which Oracle creates files in a disk group. For information about disk group types and templates, see "[Managing Disk Group Templates](#)" on page 7-15.

Oracle ASM mirroring is more flexible than traditional RAID mirroring because you can specify the redundancy level for each file. For example, two files can share the same disk group with one file being mirrored while the other is not.

When Oracle ASM allocates an extent for a mirrored file, Oracle ASM allocates a primary copy and a mirror copy. Oracle ASM chooses the disk on which to store the

mirror copy in a different failure group than the primary copy. Failure groups are used to place mirrored copies of data so that each copy is on a disk in a different failure group. The simultaneous failure of all disks in a failure group does not result in data loss.

You define the failure groups for a disk group when you create an Oracle ASM disk group. After a disk group is created, you cannot alter the redundancy level of the disk group. If you omit the failure group specification, then Oracle ASM automatically places each disk into its own failure group, except for disk groups containing disks on Oracle Exadata cells. Normal redundancy disk groups require at least two failure groups. High redundancy disk groups require at least three failure groups. Disk groups with external redundancy do not use failure groups.

For more information about mirroring and failure groups, see "[Oracle ASM Mirroring and Failure Groups](#)" on page 4-23.

About Oracle ASM Disks

Oracle ASM disks are the storage devices that are provisioned to Oracle ASM disk groups. Examples of Oracle ASM disks include:

- A disk or partition from a storage array
- An entire disk or the partitions of a disk
- Logical volumes
- Network-attached files (NFS)

When you add a disk to a disk group, you can assign an Oracle ASM disk name or Oracle ASM assigns the Oracle ASM disk name automatically. This name is different from the pathname used by the operating system. In a cluster, a disk may be assigned different operating system device names on different nodes, but the disk has the same Oracle ASM disk name on all of the nodes. In a cluster, an Oracle ASM disk must be accessible from all of the instances that share the disk group.

Oracle ASM spreads the files proportionally across all of the disks in the disk group. This allocation pattern maintains every disk at the same capacity level and ensures that all of the disks in a disk group have the same I/O load. Because Oracle ASM load balances among all of the disks in a disk group, different Oracle ASM disks should not share the same physical drive.

Allocation Units

Every Oracle ASM disk is divided into allocation units (AU). An allocation unit is the fundamental unit of allocation within a disk group. A file extent consists of one or more allocation units. An Oracle ASM file consists of one or more file extents.

When you create a disk group, you can set the Oracle ASM allocation unit size with the `AU_SIZE` disk group attribute. The values can be 1, 2, 4, 8, 16, 32, or 64 MB, depending on the specific disk group compatibility level. Larger AU sizes typically provide performance advantages for data warehouse applications that use large sequential reads. For information about disk group compatibility attributes, see "[Disk Group Compatibility](#)" on page 4-31

About Oracle ASM Files

Files that are stored in Oracle ASM disk groups are called Oracle ASM files. Each Oracle ASM file is contained within a single Oracle ASM disk group. Oracle Database communicates with Oracle ASM in terms of files. This is similar to the way Oracle

Database uses files on any file system. You can store the various file types in Oracle ASM disk groups, including:

- Control files
- Data files, temporary data files, and data file copies
- SPFILEs
- Online redo logs, archive logs, and Flashback logs
- RMAN backups
- Disaster recovery configurations
- Change tracking bitmaps
- Data Pump dumpsets

Oracle ASM automatically generates Oracle ASM file names as part of file creation and tablespace creation. Oracle ASM file names begin with a plus sign (+) followed by a disk group name. You can specify user-friendly aliases for Oracle ASM files and create a hierarchical directory structure for the aliases.

The following sections describe the Oracle ASM file components:

- [Extents](#)
- [Oracle ASM Striping](#)
- [File Templates](#)

Extents


The contents of Oracle ASM files are stored in a disk group as a set, or collection, of **extents** that are stored on individual disks within disk groups. Each extent resides on an individual disk. Extents consist of one or more allocation units (AU). To accommodate increasingly larger files, Oracle ASM uses variable size extents.

Variable size extents enable support for larger Oracle ASM data files, reduce SGA memory requirements for very large databases, and improve performance for file create and open operations. The initial extent size equals the disk group allocation unit size and it increases by a factor of 4 or 16 at predefined thresholds. This feature is automatic for newly created and resized data files when specific disk group compatibility attributes are set to 11.1 or higher. For information about compatibility attributes, see "[Disk Group Compatibility](#)" on page 4-31.

The extent size of a file varies as follows:

- Extent size is always equal to the disk group AU size for the first 20000 extent sets (0 - 19999).
- Extent size is equal to 4*AU size for the next 20000 extent sets (20000 - 39999).
- Extent size is equal to 16*AU size for the next 20000 and higher extent sets (40000+).

[Figure 1-4](#) shows the Oracle ASM file extent relationship with allocation units. The first eight extents (0 to 7) are distributed on four Oracle ASM disks and are equal to the AU size. After the first 20000 extent sets, the extent size becomes 4*AU for next 20000 extent sets (20000 - 39999). This is shown as bold rectangles labeled with the extent set numbers 20000 to 20007, and so on. The next increment for an Oracle ASM extent is 16*AU (not shown in [Figure 1-4](#)).

Figure 1–4 Oracle ASM File Allocation in a Disk Group

Oracle ASM Striping

Oracle ASM striping has two primary purposes:

- To balance loads across all of the disks in a disk group
- To reduce I/O latency

Coarse-grained striping provides load balancing for disk groups while fine-grained striping reduces latency for certain file types by spreading the load more widely.

To stripe data, Oracle ASM separates files into stripes and spreads data evenly across all of the disks in a disk group. The coarse-grained stripe size is always equal to the AU size (not the data extent size). The fine-grained stripe size always equals 128 KB in any configuration; this provides lower I/O latency for small I/O operations such as redo log writes.

File Templates

Templates are collections of attribute values that are used to specify disk regions, file mirroring, and **striping** attributes for an Oracle ASM file when it is created. When creating a file, you can include a template name and assign desired attributes based on an individual file rather than the file type.

A default template is provided for every Oracle file type, but you can customize templates to meet unique requirements. Each disk group has a default template associated with each file type.

For more information about Oracle ASM templates, see "[Managing Disk Group Templates](#)" on page 7-15.

Understanding Oracle ASM Disk Group Administration

This section describes Oracle ASM disk group administration and it contains the following topics:

- [About Discovering Disks](#)

- [About Mounting Disk Groups](#)
- [About Adding and Dropping Disks](#)
- [About Online Storage Reconfigurations and Dynamic Rebalancing](#)

About Discovering Disks

The disk discovery process locates the operating system names for disks that Oracle ASM can access. Disk discovery is also used to find all of the disks that comprise a disk group to be mounted. This can include the disks that you want to add to a disk group and the disks that you might consider adding to a disk group.

An Oracle ASM instance requires an `ASM_DISKSTRING` initialization parameter value to specify its discovery strings. Only pathnames that the Oracle ASM instance has permission to open are discovered. The exact syntax of a discovery string depends on the platform, ASMLIB libraries, and whether Oracle Exadata disks are used. The pathnames that an operating system accepts are always usable as discovery strings.

For information about the `ASM_DISKSTRING` initialization parameter, see "[ASM_DISKSTRING](#)" on page 3-7. For information about disk discovery, see "[Oracle ASM Disk Discovery](#)" on page 4-18.

About Mounting Disk Groups

A disk group must be mounted by a local Oracle ASM instance before database instances can access the files in the disk group. Mounting the disk group requires discovering all of the disks and locating the files in the disk group that is being mounted.

You can explicitly dismount a disk group. Oracle reports an error if you attempt to dismount a disk group without the `force` option when any of the disk group files are open. It is possible to have disks fail in excess of the Oracle ASM redundancy setting. If this happens, then the disk group is forcibly dismounted. If the disk group is forcibly dismounted, a database cannot access files in the disk group.

For more information about disk groups, see "[Mounting and Dismounting Disk Groups](#)" on page 4-42.

About Adding and Dropping Disks

You can add a disk to an existing disk group to add space and to improve throughput. The specified discovery string identifies the disk or disks that you want to add. The disks that you add must be discovered by every Oracle ASM instance using its `ASM_DISKSTRING` initialization parameter. After you add a disk, Oracle ASM rebalancing operations move data onto the new disk. To minimize the rebalancing I/O, it is more efficient to add multiple disks at the same time.

You can drop a disk from a disk group if it fails or to re-purpose capacity. Use the Oracle ASM disk name to drop a disk, not the discovery string device name. If an error occurs while writing to a disk, then Oracle ASM drops the disk automatically.

For more information about altering disk group membership, see "[Altering Disk Groups](#)" on page 4-11.

About Online Storage Reconfigurations and Dynamic Rebalancing

Rebalancing a disk group moves data between disks to ensure that every file is evenly spread across all of the disks in a disk group. When all of the files are evenly

dispersed, all of the disks are evenly filled to the same percentage; this ensures load balancing. Rebalancing does not relocate data based on I/O statistics nor is rebalancing started as a result of statistics. Oracle ASM rebalancing operations are controlled by the size of the disks in a disk group.

Oracle ASM automatically initiates a rebalance after storage configuration changes, such as when you add, drop, or resize disks. The power setting parameter determines the speed with which rebalancing operations occur.

You can manually start a rebalance to change the power setting of a running rebalance. A rebalance is automatically restarted if the instance on which the rebalancing is running stops. Databases can remain operational during rebalancing operations.

You can minimize the impact on database performance with the setting of the `POWER_LIMIT` initialization parameter.

For more information about the power limit setting, see "[ASM_POWER_LIMIT](#)" on page 3-8. For more information about disk rebalancing, see "[Manually Rebalancing Disk Groups](#)" on page 4-17.

Considerations for Oracle ASM Storage

This chapter discusses some points to consider about the storage subsystem before you configure Oracle Automatic Storage Management (Oracle ASM). When preparing your storage to use Oracle ASM, first determine the storage option for your system and then prepare the disk storage for your specific operating system environment.

When configuring your system's storage, you must consider the initial capacity of the system and your plans for future growth. Oracle ASM simplifies the task of accommodating growth; however, your growth plans can affect choices, such as the size of the Oracle ASM disks. You must also consider that I/O performance depends on the interconnect between the storage and host, not just the storage disks. As you scale up the number of nodes in a cluster, you also need to scale up the storage subsystem.

This chapter contains the following topics:

- [Storage Resources for Disk Groups](#)
- [Oracle ASM and Multipathing](#)
- [Recommendations for Storage Preparation](#)

Storage Resources for Disk Groups

You can create an Oracle ASM disk group using one of the following storage resources:

- Disk Partition

A disk **partition** can be the entire disk drive or a section of a disk drive. However, the Oracle ASM disk cannot be in a partition that includes the partition table because the partition table would be overwritten.

- **Logical Unit Number (LUN)**

A LUN is a disk presented to a computer system by a storage array. Oracle recommends that you use hardware RAID functionality to create LUNs. Storage hardware RAID 0+1 or RAID5, and other RAID configurations, can be provided to Oracle ASM as Oracle ASM disks.

- Logical Volume

A logical **volume** is supported in less complicated configurations where a logical volume is mapped to a LUN, or a logical volume uses disks or raw partitions. Logical volume configurations are not recommended by Oracle because they create a duplication of functionality. Oracle also does not recommend using logical **volume managers** for mirroring because Oracle ASM already provides mirroring.

- **Network File System (NFS)**

Oracle ASM supports NFS files as Oracle ASM disks. Oracle Database has built-in support for the network file system (NFS) and does not depend on operating system (OS) support for NFS. Although NFS and Oracle ASM have overlapping functionality, Oracle ASM can load balance or mirror across NFS files. Oracle ASM also supports Oracle Direct NFS (dNFS) client that integrates the NFS client functionality directly in the Oracle Database software stack.

See Also: *Oracle Grid Infrastructure Installation Guide* for your operating system for information about Oracle Direct NFS

Notes:

- Oracle ASM Dynamic Volume Manager (Oracle ADVM) does not support NFS.
 - Using Database Configuration Assistant (DBCA) or the Oracle Universal Installer (OUI) to store Oracle Clusterware or Oracle Database files directly on block or raw devices is not supported. Performing a new installation using block or raw devices is not allowed.
-
-

The procedures for preparing storage resources for Oracle ASM are:

1. Identify or create the storage devices for Oracle ASM by identifying all of the storage resource device names that you can use to create an Oracle ASM disk group. For example, on Linux systems without ASMLIB, device names are typically presented from the `/dev` directory with the `/dev/device_name_identifier` name syntax.
2. Change the ownership and the permissions on storage device resources.

For example, the following steps are required on Linux systems:

- Change the user and group ownership of devices, such as `grid:asmadmin`
For information about Oracle ASM privileges, see "[About Privileges for Oracle ASM](#)" on page 3-23.
- Change the device permissions to read/write

After you have configured Oracle ASM, ensure that disk discovery has been configured correctly by setting the `ASM_DISKSTRING` initialization parameter. For information about the `ASM_DISKSTRING` parameter, see "[ASM_DISKSTRING](#)" on page 3-7.

Note: Setting the ownership to `oracle:dba` is one example that corresponds to the default settings. A nondefault installation may require different settings. In general, the owner of the disk devices should be the same as the owner of the Oracle binary. The group ownership should be `OSDBA` of the Oracle ASM instance, which is defined at installation. For information about Oracle ASM privileges, see "[About Privileges for Oracle ASM](#)" on page 3-23.

For detailed information about preparing disks for an Oracle ASM installation, refer to your platform-specific installation guide for Oracle Database, Oracle Clusterware, and Oracle Real Application Clusters (Oracle RAC).

See Also: Oracle Exadata documentation for information about preparing Oracle Exadata storage

Oracle ASM and Multipathing

Multipathing solutions provide failover by using redundant physical path components. These components include adapters, cables, and switches that reside between the server and the storage subsystem. If one or more of these components fails, then applications can still access their data, eliminating a single point of failure with the **Storage Area Network (SAN)**, Host Bus Adapter, interface cable, or host port on a multiported storage array.

Multipathing is a software technology implemented at the operating system device driver level. Multipathing creates a pseudo device to facilitate the sharing and balancing of I/O operations across all of the available I/O paths. Multipathing also improves system performance by distributing the I/O load across all available paths, providing a higher level of data availability through automatic failover and failback.

Although Oracle ASM is not designed with multipathing functionality, Oracle ASM does operate with multipathing technologies. Multipathing technologies are available from many sources. Storage vendors offer multipathing products to support their specific storage products, while software vendors usually develop multipathing products to support several server platforms and storage products.

See Also: Your storage or software vendor multipathing documentation for more information about multipathing options for specific platforms and storage products

With Oracle ASM, you can ensure the discovery of a multipath disk by setting the value of the `ASM_DISKSTRING` initialization parameter to a pattern that matches the pseudo devices that represents the multipath disk. When I/O is sent to the pseudo device, the multipath driver intercepts it and provides load balancing to the underlying subpaths.

If Oracle ASM discovers multiple paths to the same disk device, Oracle ASM raises an error. Because a single disk can appear multiple times in a multipath configuration, you must configure Oracle ASM to discover only the multipath disk.

When using **ASMLIB** with Oracle ASM on Linux, you can ensure the discovery of the multipath disk by configuring Oracle ASM to scan the multipath disk first or to exclude the single path disks when scanning.

For information about disk discovery, see "[Oracle ASM Disk Discovery](#)" on page 4-18.

See Also:

- Support note 294869.1 at My Oracle Support (<http://metalink.oracle.com>) for information about Oracle ASM and Multipathing
- <http://www.oracle.com/technology/tech/linux/asmlib/multipath.html> for information about configuring Oracle ASMLIB with multipath disks
- Your platform-specific installation guide for information about configuring multipathing for your system

Recommendations for Storage Preparation

The following are guidelines for preparing storage for use with Oracle ASM:

- Configure two disk groups, one for the data file and the other for the fast recovery area.

See Also:

- *Oracle Database Backup and Recovery User's Guide* for information about configuring the fast recovery area
 - *Oracle Database Administrator's Guide* for information about specifying a fast recovery area
- A minimum of four LUNs (Oracle ASM disks) of equal size and performance are recommended for each disk group.
 - Ensure that all Oracle ASM disks in a disk group use have similar storage performance and availability characteristics. In storage configurations with mixed speed drives, such as 10K and 15K RPM, I/O performance is constrained by the slowest speed drive.
 - Oracle ASM data distribution policy is capacity-based. Ensure that Oracle ASM disks in a disk group have the same capacity to maintain balance.
 - Create external redundancy disk groups when using high-end storage arrays. High-end storage arrays generally provide hardware RAID protection. Use Oracle ASM mirroring redundancy in the absence of a hardware RAID, or when you need host-based volume management functionality, such as mirroring across storage systems. You can use Oracle ASM mirroring in configurations when mirroring between geographically-separated sites (extended clusters).
 - Minimize I/O contention between Oracle ASM disks and other applications by dedicating disks in Oracle ASM disk groups.
 - Choose a hardware RAID stripe size that is a power of 2 and less than or equal to the size of the Oracle ASM allocation unit.
 - For Linux, use the Oracle ASMLIB feature to provide consistent device naming and permission persistency.

See Also:

- The Oracle ASM page on the Oracle Technology Network Web site at <http://www.oracle.com/technology/products/databases/asm> for more information about ASMLIB
- The Oracle ASMLIB page on the Oracle Technology Network Web site at <http://www.oracle.com/technology/tech/linux/asmlib> to download ASMLIB

Administering Oracle ASM Instances

This chapter describes how to administer Automatic Storage Management (Oracle ASM) instances. It explains how to configure Oracle ASM instance parameters as well how to set Oracle Database parameters for use with Oracle ASM. The chapter also describes Oracle ASM upgrading, patching, and authentication for Oracle ASM instance access. You can also use procedures in this chapter to migrate a database to use Oracle ASM.

Administering an Oracle ASM instance is similar to administering an Oracle Database instance, but the process requires fewer procedures. You can use Oracle Enterprise Manager and SQL*Plus to perform Oracle ASM instance administration tasks.

Oracle ASM is installed in the Oracle grid infrastructure home separate from the Oracle Database home. When managing an Oracle ASM instance, the administration activity must be performed in the Oracle grid infrastructure home.

This chapter contains the following topics:

- [Operating With Different Releases of Oracle ASM and Database Instances Simultaneously](#)
- [Configuring Initialization Parameters for an Oracle ASM Instance](#)
- [Managing Oracle ASM Instances](#)
- [Using Oracle ASM Rolling Upgrade](#)
- [Patching Oracle ASM Instances](#)
- [Authentication for Accessing Oracle ASM Instances](#)
- [Migrating a Database to Use Oracle ASM](#)

For a description of an Oracle ASM instance, see "[About Oracle ASM Instances](#)" on page 1-2. For information about using Oracle Enterprise Manager to administer Oracle ASM, see [Chapter 9, "Administering Oracle ASM with Oracle Enterprise Manager"](#).

Operating With Different Releases of Oracle ASM and Database Instances Simultaneously

Oracle Automatic Storage Management (Oracle ASM) in Oracle Database 11g Release 2 (11.2) supports 11g Release 2 (11.2) or older software versions of Oracle database instances, including Oracle Database 10g. For compatibility between Oracle Clusterware and Oracle ASM, the Oracle Clusterware release must be greater than or equal to the Oracle ASM release.

Notes:

- An Oracle ASM instance must be at 11g Release 2 (11.2) to support an 11g Release 2 (11.2) Oracle Database.
 - See *Oracle Exadata documentation* for information about the Oracle Database versions that Oracle ASM 11g Release 2 (11.2) supports when Oracle Exadata storage is present.
-
-

There are additional compatibility considerations when using disk groups with different releases of Oracle ASM and database instances. For information about disk group compatibility attributes settings, see "[Disk Group Compatibility](#)" on page 4-31.

When using different software versions, the database instance supports Oracle ASM functionality of the earliest release in use. For example, a 10.1 database instance operating with an 11.2 Oracle ASM instance supports only Oracle ASM 10.1 features.

The `V$ASM_CLIENT` view contains the `SOFTWARE_VERSION` and `COMPATIBLE_VERSION` columns with information about the software version number and instance compatibility level.

- The `SOFTWARE_VERSION` column of `V$ASM_CLIENT` contains the software version number of the database or Oracle ASM instance for the selected disk group connection.
- The `COMPATIBLE_VERSION` column contains the setting of `COMPATIBLE` parameter of the database or Oracle ASM instance for the selected disk group connection.

You can query the `V$ASM_CLIENT` view on both Oracle ASM and database instances. For an example showing a query on the `V$ASM_CLIENT` view, see [Example 6-4, "Viewing Disk Group Clients With V\\$ASM_CLIENT"](#) on page 6-3. For more information about the `V$ASM_CLIENT` and `V$ASM_*` views, see "[Views Containing Oracle ASM Disk Group Information](#)" on page 6-1.

Configuring Initialization Parameters for an Oracle ASM Instance

This section discusses initialization parameter files and parameter settings for Oracle ASM instances. To install and initially configure an Oracle ASM instance, use Oracle Universal Installer (OUI) and Oracle ASM Configuration Assistant (ASMCA). Refer to your platform-specific *Oracle Grid Infrastructure Installation Guide* for details about installing and configuring Oracle ASM.

After an Oracle ASM instance has been installed on a single-instance Oracle Database or in an Oracle Real Application Clusters (Oracle RAC) environment, the final Oracle ASM configuration can be performed. You only need to configure a few Oracle ASM-specific instance initialization parameters. The default values are sufficient in most cases.

See Also: <http://www.oracle.com/technology/asm/> for more information about Oracle ASM best practices

This section contains the following topics:

- [Initialization Parameter Files for an Oracle ASM Instance](#)
- [Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File](#)
- [Setting Oracle ASM Initialization Parameters](#)

- [Oracle ASM Parameter Setting Recommendations](#)
- [Setting Database Initialization Parameters for Use with Oracle ASM](#)

See Also:

- *Oracle Database Reference* for information about initialization parameters
- *Oracle Database Administrator's Guide* for information about initialization parameter files

Initialization Parameter Files for an Oracle ASM Instance

When installing Oracle ASM in an Oracle Restart configuration, Oracle Universal Installer (OUI) creates a separate server parameter file (SPFILE) and password file for the Oracle ASM instance. When installing Oracle ASM in a clustered Oracle ASM environment where the Oracle grid infrastructure home is shared among all of the nodes, OUI creates a single SPFILE for Oracle ASM in a disk group. In a clustered environment without a shared Oracle grid infrastructure home, OUI creates a SPFILE for Oracle ASM on each node.

You can use an SPFILE or a text-based initialization parameter file (PFILE) as the Oracle ASM instance parameter file. If you use an SPFILE in a clustered Oracle ASM environment, then you must place the SPFILE in a disk group, a shared raw device, or on a cluster file system. Oracle recommends that the Oracle ASM SPFILE is placed in a disk group. You cannot use a new alias created on an existing Oracle ASM SPFILE to start up the Oracle ASM instance

If you do not use a shared Oracle grid infrastructure home, then the Oracle ASM instance can use a PFILE. The same rules for file name, default location, and search order that apply to database initialization parameter files also apply to Oracle ASM initialization parameter files.

When an Oracle ASM instance searches for an initialization parameter file, the search order is:

1. The location of the initialization parameter file specified in the Grid Plug and Play (GPnP) profile
2. If the location has not been set in the GPnP profile, the search order changes to:
 - a. SPFILE in the Oracle ASM instance home

For example, the SPFILE for Oracle ASM has the following default path in the Oracle grid infrastructure home in a Linux environment:

```
$ORACLE_HOME/dbs/spfile+ASM.ora
```

- b. PFILE in the Oracle ASM instance home

You can administer an Oracle ASM initialization parameter files with SQL*Plus, Oracle Enterprise Manager, ASMCA, and ASMCMD commands. For information about Oracle Enterprise Manager, see "[Configuring Oracle ASM Initialization Parameters with Oracle Enterprise Manager](#)" on page 9-3. For information about ASMCA, see [Chapter 11, "Oracle ASM Configuration Assistant"](#). For information about ASMCMD commands, see "[ASMCMD Instance Management Commands](#)" on page 12-9.

See Also:

- *Oracle Database Administrator's Guide* for more information about creating and maintaining an initialization parameter files
- *Oracle Database 2 Day DBA* for information about viewing and modifying initialization parameters
- *Oracle Database SQL Language Reference* for information about creating an SPFILE with the `CREATE SPFILE` SQL statement

Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File

You can back up, copy, or move an Oracle ASM SPFILE with the `ASMCMD spbackup`, `spcopy` or `spmmove` commands. For information about these `ASMCMD` commands, see "[spbackup](#)" on page 12-14, "[spcopy](#)" on page 12-15, and "[spmmove](#)" on page 12-16.

You can also use the SQL `CREATE SPFILE` to create an Oracle ASM SPFILE when connected to the Oracle ASM instance.

You can copy and move an Oracle ASM PFILE with the commands available on the specific platform, such as `cp` for Linux.

After copying or moving a SPFILE or PFILE, you must restart the instance with the SPFILE or PFILE in the new location if you want to use that SPFILE or PFILE.

If the `COMPATIBLE.ASM` disk group attribute is set to 11.2 or greater for a disk group, you can create, copy, or move an Oracle ASM SPFILE into the disk group.

For example, after upgrading an Oracle ASM instance from 11g release 1 (11.1) to 11g release 2 (11.2), you could place the Oracle ASM SPFILE in a disk group that has `COMPATIBLE.ASM` set to 11.2.

In the following steps, assume an Oracle ASM 11g release 2 (11.2) instance is using a PFILE stored in `$ORACLE_HOME/dbs/asmspfile.ora`. You can use the SQL `CREATE SPFILE` statement to create a SPFILE from a PFILE stored in a local or shared file system. If a PFILE does not exist, then it could be created with the SQL `CREATE PFILE` statement.

To create a SPFILE in a disk group, perform the following steps:

1. Connect to the Oracle ASM instance.

For example:

```
$ sqlplus / as sysasm
```

2. Create a SPFILE in a disk group that has `COMPATIBLE.ASM` set to 11.2 with the SQL `CREATE SPFILE` statement.

For example, create an Oracle ASM SPFILE from the existing PFILE.

```
SQL> CREATE SPFILE = '+DATA/asmspfile.ora'
      FROM PFILE = '$ORACLE_HOME/dbs/asmspfile.ora';
```

The `CREATE SPFILE` statement also updates the Grid Plug and Play (GPnP) profile. You can check the location of the Oracle ASM SPFILE in the GPnP profile with the `ASMCMD spget` command. See "[spget](#)" on page 12-16.

3. Restart the Oracle ASM instance so that the instance will use the SPFILE in the new location.

For information on shutting down and starting up an Oracle ASM instance, see ["Starting Up an Oracle ASM Instance"](#) on page 3-13 and ["Shutting Down an Oracle ASM Instance"](#) on page 3-16.

For information about disk group compatibility attributes, see ["Disk Group Compatibility"](#) on page 4-31. For information about upgrading an Oracle ASM instance, see ["Upgrading an Oracle ASM Instance With Oracle Universal Installer"](#) on page 3-17.

See Also:

- *Oracle Database Administrator's Guide* for more information about creating and maintaining an initialization parameter files
- *Oracle Database SQL Language Reference* for information about creating an SPFILE with the `CREATE SPFILE SQL` statement
- *Oracle Real Application Clusters Installation Guide* for information about Grid Plug and Play (GPnP)

Setting Oracle ASM Initialization Parameters

There are several initialization parameters that you must set for an Oracle ASM instance. You can set these parameters with Oracle ASM Configuration Assistant (ASMCA). You can also set some of these parameters after database creation using Oracle Enterprise Manager or SQL `ALTER SYSTEM` or `ALTER SESSION` statements.

The `INSTANCE_TYPE` initialization parameter is the only required parameter in the Oracle ASM instance parameter file. The Oracle ASM* parameters use suitable defaults for most environments. You cannot use parameters with names that are prefixed with Oracle ASM* in database instance parameter files.

Some database initialization parameters are also valid for an Oracle ASM instance initialization file. In general, Oracle ASM selects the appropriate defaults for database parameters that are relevant to an Oracle ASM instance.

For information about setting Oracle ASM parameters with Oracle Enterprise Manager, see ["Configuring Oracle ASM Initialization Parameters with Oracle Enterprise Manager"](#) on page 9-3.

Automatic Memory Management for Oracle ASM

Automatic memory management automatically manages the memory-related parameters for both Oracle ASM and database instances with the `MEMORY_TARGET` parameter. Automatic memory management is enabled by default on an Oracle ASM instance, even when the `MEMORY_TARGET` parameter is not explicitly set. The default value used for `MEMORY_TARGET` is acceptable for most environments. This is the only parameter that you must set for complete Oracle ASM memory management. Oracle strongly recommends that you use automatic memory management for Oracle ASM.

If you do not set a value for `MEMORY_TARGET`, but you do set values for other memory related parameters, Oracle internally calculates the optimum value for `MEMORY_TARGET` based on those memory parameter values. You can also increase `MEMORY_TARGET` dynamically, up to the value of the `MEMORY_MAX_TARGET` parameter, just as you can do for the database instance.

Although it is not recommended, you can disable automatic memory management by either setting the value for `MEMORY_TARGET` to 0 in the Oracle ASM parameter file or by running an `ALTER SYSTEM SET MEMORY_TARGET=0` statement. When you disable automatic memory management, Oracle reverts to auto shared memory management and automatic PGA memory management. If you want to revert to Oracle Database

10g release 2 (10.2) functionality to manually manage Oracle ASM SGA memory, also run the `ALTER SYSTEM SET SGA_TARGET=0` statement. You can then manually manage Oracle ASM memory using the information in "[Oracle ASM Parameter Setting Recommendations](#)" on page 3-6, that discusses Oracle ASM memory-based parameter settings. Unless specified, the behaviors of all of the automatic memory management parameters in Oracle ASM instances is the same as in Oracle Database instances.

Notes:

- For a Linux environment, automatic memory management cannot work if `/dev/shm` is not available or is undersized. For more information, see *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems*. For information about platforms that support automatic memory management, see *Oracle Database Administrator's Guide*.
 - The minimum `MEMORY_TARGET` for Oracle ASM is 256 MB. If you set `MEMORY_TARGET` to 100 MB, then Oracle increases the value for `MEMORY_TARGET` to 256 MB automatically.
-

See Also:

- *Oracle Database Administrator's Guide* for more information about the functionality of automatic memory management for database instances, which varies from Oracle ASM
- *Oracle Database Concepts* for an overview of memory management methods

Oracle ASM Parameter Setting Recommendations

This section contains information about the following parameters for Oracle ASM:

- [ASM_DISKGROUPS](#)
- [ASM_DISKSTRING](#)
- [ASM_POWER_LIMIT](#)
- [ASM_PREFERRED_READ_FAILURE_GROUPS](#)
- [DB_CACHE_SIZE](#)
- [DIAGNOSTIC_DEST](#)
- [INSTANCE_TYPE](#)
- [LARGE_POOL_SIZE](#)
- [PROCESSES](#)
- [REMOTE_LOGIN_PASSWORDFILE](#)
- [SHARED_POOL_SIZE](#)

See Also:

- *Oracle Database Administrator's Guide* for more information about creating and maintaining an initialization parameter file
- *Oracle Database 2 Day DBA* for information about viewing and modifying initialization parameters

ASM_DISKGROUPS

The `ASM_DISKGROUPS` initialization parameter specifies a list of the names of disk groups that an Oracle ASM instance mounts at startup. Oracle ignores the value that you set for `ASM_DISKGROUPS` when you specify the `NOMOUNT` option at startup or when you issue the `ALTER DISKGROUP ALL MOUNT` statement. The default value of the `ASM_DISKGROUPS` parameter is a `NULL` string. For information about disk groups that are mounted at startup time, see ["About Mounting Disk Groups at Startup"](#) on page 3-15.

The `ASM_DISKGROUPS` parameter is dynamic. If you are using a server parameter file (SPFILE), then you should not need to manually alter the value of `ASM_DISKGROUPS`. Oracle ASM automatically adds a disk group to this parameter when the disk group is successfully created or mounted. Oracle ASM also automatically removes a disk group from this parameter when the disk group is dropped or dismounted.

The following is an example of setting the `ASM_DISKGROUPS` parameter dynamically:

```
SQL> ALTER SYSTEM SET ASM_DISKGROUPS = DATA, FRA;
```

When using a text initialization parameter file (PFILE), you may edit the initialization parameter file to add the name of any disk group so that it is mounted automatically at instance startup. You must remove the name of any disk group that you no longer want automatically mounted.

The following is an example of the `ASM_DISKGROUPS` parameter in the initialization file:

```
ASM_DISKGROUPS = DATA, FRA
```

Note: Issuing the `ALTER DISKGROUP...ALL MOUNT` or `ALTER DISKGROUP...ALL DISMOUNT` commands does not affect the value of `ASM_DISKGROUPS`.

For additional information about mounting Oracle ASM disk groups, see ["Mounting and Dismounting Disk Groups"](#) on page 4-42.

See Also: *Oracle Database Reference* for more information about the `ASM_DISKGROUPS` initialization parameter

ASM_DISKSTRING

The `ASM_DISKSTRING` initialization parameter specifies a comma-delimited list of strings that limits the set of disks that an Oracle ASM instance discovers. The discovery strings can include wildcard characters. Only disks that match one of the strings are discovered. The same disk cannot be discovered twice.

The discovery string format depends on the Oracle ASM library and the operating system that are in use. Pattern matching is supported. Refer to your operating system-specific installation guide for information about the default pattern matching.

For example, on a Linux server that does not use ASMLIB, to limit the discovery process to only include disks that are in the `/dev/rdisk/mydisks` directory, set the `ASM_DISKSTRING` initialization parameter to:

```
/dev/rdisk/mydisks/*
```

The asterisk is required. To limit the discovery process to only include disks that have a name that ends in `disk3` or `disk4`, set `ASM_DISKSTRING` to:

```
/dev/rdisk/*disk3, /dev/rdisk/*disk4
```

The ? character, when used as the first character of a path, expands to the Oracle home directory. Depending on the operating system, when you use the ? character elsewhere in the path, it is a wildcard for one character.

The default value of the `ASM_DISKSTRING` parameter is a NULL string. A NULL value causes Oracle ASM to search a default path for all disks in the system to which the Oracle ASM instance has read and write access. The default search path is platform-specific. Refer to your operating system specific installation guide for more information about the default search path.

Oracle ASM cannot use a disk unless all of the Oracle ASM instances in the cluster can discover the disk through one of their own discovery strings. The names do not need to be the same on every node, but all disks must be discoverable by all of the nodes in the cluster. This may require dynamically changing the initialization parameter to enable adding new storage.

For additional information about discovering disks, see "[Oracle ASM Disk Discovery](#)" on page 4-18.

See Also:

- *Oracle Exadata documentation* for information about the Oracle ASM discovery string format for Oracle Exadata
- *Oracle Database Reference* for more information about the `ASM_DISKSTRING` initialization parameter

ASM_POWER_LIMIT

The `ASM_POWER_LIMIT` initialization parameter specifies the default power for disk rebalancing. The default value is 1 and the range of allowable values is 0 to 11 inclusive. A value of 0 disables rebalancing. Higher numeric values enable the rebalancing operation to complete more quickly, but might result in higher I/O overhead.

See Also: "[Tuning Rebalance Operations](#)" on page 4-18 for more information about `ASM_POWER_LIMIT` and *Oracle Database Reference* for more information about the `ASM_POWER_LIMIT` initialization parameter

ASM_PREFERRED_READ_FAILURE_GROUPS

The `ASM_PREFERRED_READ_FAILURE_GROUPS` initialization parameter value is a comma-delimited list of strings that specifies the failure groups that should be preferentially read by the given instance. This parameter is generally used only for clustered Oracle ASM instances and its value can be different on different nodes. For example:

diskgroup_name1.failure_group_name1, ...

The `ASM_PREFERRED_READ_FAILURE_GROUPS` parameter setting is instance specific. This parameter is only valid for clustered Oracle ASM instances and the default value is NULL.

Note: The `ASM_PREFERRED_READ_FAILURE_GROUPS` parameter is valid only in Oracle RAC environments.

See Also:

- ["Preferred Read Failure Groups"](#) on page 4-29 for more information about `ASM_PREFERRED_READ_FAILURE_GROUPS`
- *Oracle Real Application Clusters Administration and Deployment Guide* for more information about configuring preferred disks in extended clusters
- *Oracle Database Reference* for more information about the `ASM_PREFERRED_READ_FAILURE_DISKGROUPS` initialization parameter

DB_CACHE_SIZE

You do not need to set a value for the `DB_CACHE_SIZE` initialization parameter if you use automatic memory management.

The setting for the `DB_CACHE_SIZE` parameter determines the size of the buffer cache. This buffer cache is used to store metadata blocks. The default value for this parameter is suitable for most environments.

See Also:

- *Oracle Database Administrator's Guide* for more information about setting the `DB_CACHE_SIZE` initialization parameter
- *Oracle Database Performance Tuning Guide* for more information about memory configuration
- *Oracle Database Reference* for more information about the `DB_CACHE_SIZE` parameter

DIAGNOSTIC_DEST

The `DIAGNOSTIC_DEST` initialization parameter specifies the directory where diagnostics for an instance are located. The default value for an Oracle ASM instance is the `$ORACLE_BASE` directory.

See Also:

- *Oracle Database Administrator's Guide* for more information about setting the `DIAGNOSTIC_DEST` initialization parameter
- *Oracle Database Reference* for more information about the `DIAGNOSTIC_DEST` parameter

INSTANCE_TYPE

The `INSTANCE_TYPE` initialization parameter must be set to `Oracle ASM` for an Oracle ASM instance. This parameter is optional for an Oracle ASM instance in an Oracle grid infrastructure home.

The following is an example of the `INSTANCE_TYPE` parameter in the initialization file:

```
INSTANCE_TYPE = ASM
```

See Also: *Oracle Database Reference* for more information about the `INSTANCE_TYPE` parameter

LARGE_POOL_SIZE

You do not need to set a value for the `LARGE_POOL_SIZE` initialization parameter if you use automatic memory management.

The setting for the `LARGE_POOL_SIZE` parameter is used for large allocations. The default value for this parameter is suitable for most environments.

See Also:

- *Oracle Database Administrator's Guide* for more information about setting the `LARGE_POOL_SIZE` initialization parameter
- *Oracle Database Performance Tuning Guide* for more information about memory configuration
- *Oracle Database Reference* for more information about the `LARGE_POOL_SIZE` parameter

PROCESSES

The `PROCESSES` initialization parameter affects Oracle ASM, but generally you do not need to modify the setting. The default value provided is usually suitable.

See Also:

- *Oracle Database Administrator's Guide* for more information about setting the `PROCESSES` initialization parameter
- *Oracle Database Reference* for more information about the `PROCESSES` parameter

REMOTE_LOGIN_PASSWORDFILE

The `REMOTE_LOGIN_PASSWORDFILE` initialization parameter specifies whether the Oracle ASM instance checks for a password file. This parameter operates the same for Oracle ASM and database instances.

See Also:

- *Oracle Database Administrator's Guide* for more information about setting the `REMOTE_LOGIN_PASSWORDFILE` initialization parameter
- *Oracle Database Reference* for more information about the `REMOTE_LOGIN_PASSWORDFILE` parameter

SHARED_POOL_SIZE

You do not need to set a value for the `SHARED_POOL_SIZE` initialization parameter if you use automatic memory management.

The setting for the `SHARED_POOL_SIZE` parameter determines the amount of memory required to manage the instance. The setting for this parameter is also used to determine the amount of space that is allocated for extent storage. The default value for this parameter is suitable for most environments.

See Also:

- *Oracle Database Administrator's Guide* for more information about setting the SHARED_POOL_SIZE initialization parameter
- *Oracle Database Performance Tuning Guide* for more information about memory configuration
- *Oracle Database Reference* for more information about the SHARED_POOL_SIZE parameter

Setting Database Initialization Parameters for Use with Oracle ASM

When you do not use automatic memory management in a database instance, the SGA parameter settings for a database instance may require minor modifications to support Oracle ASM. When you use automatic memory management, the sizing data discussed in this section can be treated as informational only or as supplemental information to help determine the appropriate values that you should use for the SGA. Oracle highly recommends using automatic memory management.

See Also:

- *Oracle Database Administrator's Guide* for information about managing memory allocation in an Oracle Database instance
- *Oracle Database Performance Tuning Guide* for more information about memory configuration and use

The following are guidelines for SGA sizing on the database instance:

- PROCESSES initialization parameter—Add 16 to the current value
- LARGE_POOL_SIZE initialization parameter—Add an additional 600K to the current value
- SHARED_POOL_SIZE initialization parameter—Aggregate the values from the following queries to obtain the current database storage size that is either already on Oracle ASM or will be stored in Oracle ASM. Next, determine the redundancy type and calculate the SHARED_POOL_SIZE using the aggregated value as input.

```
SELECT SUM(bytes)/(1024*1024*1024) FROM V$DATAFILE;
SELECT SUM(bytes)/(1024*1024*1024) FROM V$LOGFILE a, V$LOG b
WHERE a.group#=b.group#;
SELECT SUM(bytes)/(1024*1024*1024) FROM V$TEMPFILE
WHERE status='ONLINE';
```

- For disk groups using external redundancy, every 100 GB of space needs 1 MB of extra shared pool plus 2 MB
- For disk groups using normal redundancy, every 50 GB of space needs 1 MB of extra shared pool plus 4 MB
- For disk groups using high redundancy, every 33 GB of space needs 1 MB of extra shared pool plus 6 MB

See Also:

- *Oracle Database Administrator's Guide* for information about managing memory allocation in an Oracle Database instance
- *Oracle Database Performance Tuning Guide* for more information about memory configuration and use

Managing Oracle ASM Instances

The following section describes how to administer Oracle ASM instances under the following topics:

- [Administering Oracle ASM Instances with Server Control Utility](#)
- [Using Oracle Restart](#)
- [Starting Up an Oracle ASM Instance](#)
- [Shutting Down an Oracle ASM Instance](#)
- [Upgrading an Oracle ASM Instance With Oracle Universal Installer](#)
- [Downgrading an Oracle ASM Instance](#)
- [Active Session History Sampling for Oracle ASM](#)

Administering Oracle ASM Instances with Server Control Utility

In addition to the Oracle ASM administration procedures that this section describes, you can use Server Control Utility (SRVCTL) in clustered Oracle ASM environments to perform the following Oracle ASM administration tasks:

- Add and remove Oracle ASM instance records in the Oracle Cluster Registry (OCR)
- Enable, disable, start, and stop Oracle ASM instance
- Display the Oracle ASM instance configuration and status

See Also: The *Oracle Real Application Clusters Administration and Deployment Guide* for information about administering Oracle ASM instances with SRVCTL

Using Oracle Restart

Oracle Restart improves the availability of your Oracle database. When you install the Oracle grid infrastructure for a standalone server, it includes both Oracle ASM and Oracle Restart. Oracle Restart runs out of the Oracle grid infrastructure home, which you install separately from Oracle Database homes.

Oracle Restart provides managed startup and restart of a single-instance (non-clustered) Oracle Database, Oracle ASM instance, service, listener, and any other process running on the server. If an interruption of a service occurs after a hardware or software failure, Oracle Restart automatically takes the necessary steps to restart the component.

With Server Control Utility (SRVCTL) you can add a component, such as an Oracle ASM instance, to Oracle Restart. You then enable Oracle Restart protection for the Oracle ASM instance. With SRVCTL, you also remove or disable Oracle Restart protection.

See Also:

- *Oracle Database Administrator's Guide* for information about configuring and administering Oracle Restart
- *Oracle Real Application Clusters Administration and Deployment Guide* for information about automatically restarting single-instance databases residing on an Oracle RAC node
- *Oracle Grid Infrastructure Installation Guide* for information about installation of Oracle grid infrastructure

Starting Up an Oracle ASM Instance

You start an Oracle ASM instance similarly to the way in which you start an Oracle database instance with some minor differences. When starting an Oracle ASM instance, note the following:

- To connect to a local Oracle ASM instance with SQL*Plus, set the `ORACLE_SID` environment variable to the Oracle ASM SID.

The default Oracle ASM SID for a single-instance database is `+ASM`, and the default SID for Oracle ASM for an Oracle RAC node is `+ASMnode_number` where `node_number` is the number of the node. The `ORACLE_HOME` environment variable must be set the Grid Infrastructure home where Oracle ASM was installed.

- The initialization parameter file must contain the following entry:

```
INSTANCE_TYPE = ASM
```

This parameter indicates that an Oracle ASM instance, not a database instance, is starting.

- When you run the `STARTUP` command, rather than trying to mount and open a database, this command attempts to mount Oracle ASM disk groups.

For information about disk groups that are mounted at startup time, see ["About Mounting Disk Groups at Startup"](#) on page 3-15.

After the Oracle ASM instance has started, you can mount disk groups with the `ALTER DISKGROUP . . . MOUNT` command. See ["Mounting and Dismounting Disk Groups"](#) on page 4-42 for more information.

- The associated Oracle database instance does not need to be running when you start the associated Oracle ASM instance.

The following list describes how Oracle ASM interprets SQL*Plus `STARTUP` command parameters.

- **FORCE** Parameter

Issues a `SHUTDOWN ABORT` to the Oracle ASM instance before restarting it.

If an Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file system is currently mounted on Oracle ADVM volumes, the file system should first be dismounted. Otherwise, applications will encounter I/O errors and Oracle ACFS user data and metadata may not be flushed to storage before the Oracle ASM storage is fenced. For information about dismounting an Oracle ACFS file system, see ["Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems"](#) on page 13-5.

- **MOUNT** or **OPEN** Parameter

Mounts the disk groups specified in the `ASM_DISKGROUPS` initialization parameter. This is the default if no command parameter is specified.

- **NOMOUNT** Parameter
Starts up the Oracle ASM instance without mounting any disk groups.
- **RESTRICT** Parameter
Starts up an instance in restricted mode that enables access only to users with both the `CREATE SESSION` and `RESTRICTED SESSION` system privileges. The `RESTRICT` clause can be used in combination with the `MOUNT`, `NOMOUNT`, and `OPEN` clauses.

See Also: ["About Restricted Mode"](#) on page 3-15 for more information

In restricted mode, database instances cannot use the disk groups. In other words, databases cannot open files that are in that disk group. Also, the disk group cannot be mounted by any other instance in the cluster. Mounting the disk group in restricted mode enables only one Oracle ASM instance to mount the disk group. This mode is useful to mount the disk group for repairing configuration issues.

The following is a sample SQL*Plus session for starting an Oracle ASM instance.

```
SQLPLUS /NOLOG
SQL> CONNECT SYS AS SYSASM
Enter password: sys_password
Connected to an idle instance.

SQL> STARTUP
ASM instance started

Total System Global Area      71303168 bytes
Fixed Size                     1069292 bytes
Variable Size                  45068052 bytes
ASM Cache                      25165824 bytes
ASM disk groups mounted
```

For more information about user authentication, see ["Authentication for Accessing Oracle ASM Instances"](#) on page 3-22.

See Also:

- *Oracle Database Administrator's Guide* for more information about using environment variables to select instances
- *Oracle Database Administrator's Guide* for more information about starting up and shutting down Oracle instances
- *Oracle Real Application Clusters Administration and Deployment Guide* for information about starting an Oracle ASM instance with SRVCTL in Oracle RAC
- *Oracle Clusterware Administration and Deployment Guide* for information about Oracle Clusterware Cluster subcomponent processes and background processes
- *Oracle Database Concepts* for information about Oracle database background processes
- *Oracle Database Reference* for a description of the Oracle background processes

About Mounting Disk Groups at Startup

At startup, the Oracle ASM instance attempts to mount the following disk groups:

- Disk groups specified in the `ASM_DISKGROUPS` initialization parameter
- Disk group used by Cluster Synchronization Services (CSS) for voting files
- Disk groups used by Oracle Clusterware for Oracle Cluster Registry (OCR)
- Disk group used by the Oracle ASM instance to store the ASM server parameter file (SPFILE)

If no disk groups are found in the previous list, then the Oracle ASM instance does not mount any disk groups at startup.

After the Oracle ASM instance has started, you can mount disk groups with the `ALTER DISKGROUP . . . MOUNT` command. For more information, see "[Mounting and Dismounting Disk Groups](#)" on page 4-42.

About Restricted Mode

You can use the `STARTUP RESTRICT` command to control access to an Oracle ASM instance while you perform maintenance. When an Oracle ASM instance is active in this mode, all of the disk groups that are defined in the `ASM_DISKGROUPS` parameter are mounted in `RESTRICTED` mode. This prevents databases from connecting to the Oracle ASM instance. In addition, the `restricted` clause of the `ALTER SYSTEM` statement is disabled for the Oracle ASM instance. The `ALTER DISKGROUP diskgroup MOUNT` statement is extended to enable Oracle ASM to mount a disk group in restricted mode.

When you mount a disk group in `RESTRICTED` mode, the disk group can only be mounted by one instance. Clients of Oracle ASM on that node cannot access that disk group while the disk group is mounted in `RESTRICTED` mode. The `RESTRICTED` mode enables you to perform maintenance tasks on a disk group in the Oracle ASM instance without interference from clients.

Rebalance operations that occur while a disk group is in `RESTRICTED` mode eliminate the lock and unlock extent map messaging that occurs between Oracle ASM instances in an Oracle RAC environment. This improves the overall rebalance throughput. At the end of a maintenance period, you must explicitly dismount the disk group and remount it in normal mode.

Shutting Down an Oracle ASM Instance

The Oracle ASM shutdown process is initiated when you run the `SHUTDOWN` command in SQL*Plus. Before you run this command, ensure that the `ORACLE_SID` environment variable is set to the Oracle ASM SID so that you can connect to the local Oracle ASM instance. The default Oracle ASM SID for a single-instance database is `+ASM`, and the default SID for Oracle ASM for an Oracle RAC node is `+ASMnode_number` where `node_number` is the number of the node. The `ORACLE_HOME` environment variable must be set the Grid Infrastructure home where Oracle ASM was installed.

Oracle strongly recommends that you shut down all database instances that use the Oracle ASM instance and dismount all file systems mounted on Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes before attempting to shut down the Oracle ASM instance.

See Also:

- *Oracle Database Administrator's Guide* for more information about using environment variables to select instances
- *Oracle Database Administrator's Guide* for more information about starting up and shutting down Oracle instances
- *Oracle Clusterware Administration and Deployment Guide* for information about shutting down an Oracle ASM instance when voting files are stored in a disk group.

```
SQLPLUS /NOLOG
SQL> CONNECT SYS AS SYSASM
Enter password: sys_password
Connected.
SQL> SHUTDOWN NORMAL
```

For more information about user authentication, see "[Authentication for Accessing Oracle ASM Instances](#)" on page 3-22.

The following list describes the `SHUTDOWN` modes and describes the behavior of the Oracle ASM instance in each mode.

- `NORMAL` Clause
Oracle ASM waits for any in-progress SQL to complete before performing an orderly dismount of all of the disk groups and shutting down the Oracle ASM instance. Before the instance is shut down, Oracle ASM waits for all of the currently connected users to disconnect from the instance. If any database instances are connected to the Oracle ASM instance, then the `SHUTDOWN` command returns an error and leaves the Oracle ASM instance running. `NORMAL` is the default shutdown mode.
- `IMMEDIATE` or `TRANSACTIONAL` Clause
Oracle ASM waits for any in-progress SQL to complete before performing an orderly dismount of all of the disk groups and shutting down the Oracle ASM instance. Oracle ASM does not wait for users currently connected to the instance to disconnect. If any database instances are connected to the Oracle ASM instance, then the `SHUTDOWN` command returns an error and leaves the Oracle ASM instance running. Because the Oracle ASM instance does not contain any transactions, the `TRANSACTIONAL` mode is the same as the `IMMEDIATE` mode.
- `ABORT` Clause

The Oracle ASM instance immediately shuts down without the orderly dismount of disk groups. This causes recovery to occur upon the next Oracle ASM startup.

If any database instance is connected to the Oracle ASM instance, then the database instance aborts.

If any Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems are currently mounted on Oracle ADVM volumes, those file systems should first be dismounted. Otherwise, applications will encounter I/O errors and Oracle ACFS user data and metadata may not be flushed to storage before the Oracle ASM storage is fenced. For information about dismounting an Oracle ACFS file system, see ["Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems"](#) on page 13-5. For more information about user authentication on Oracle ASM instance, see ["Authentication for Accessing Oracle ASM Instances"](#) on page 3-22.

Upgrading an Oracle ASM Instance With Oracle Universal Installer

This section discusses the process to upgrade an Oracle ASM instance to an Oracle Restart 11g release 2 (11.2) configuration. The recommended practice is to upgrade an Oracle ASM instance with Oracle grid infrastructure Oracle Universal Installer (OUI). OUI automatically defaults to upgrade mode when it detects an Oracle ASM instance at a previous release level. Before you make any changes to the Oracle software, Oracle recommends that you create a backup of the Oracle software.

The following procedure describes how to upgrade an Oracle ASM instance from 11g release 1 (11.1) to 11g release 2 (11.2) in an Oracle Restart configuration. In this scenario:

- The Oracle ASM and Oracle Database instances 11g release 1 (11.1) exist in separate homes.
 - The Oracle grid infrastructure 11g release 2 (11.2) is to be installed in a separate home and the Oracle ASM instance 11g release 2 (11.2) is to be set up as an Oracle Restart (single-instance) configuration.
1. Shut down the Oracle Enterprise Manager agent, Oracle Database instances, Oracle ASM instance, and the listener the in the older database and Oracle ASM homes.
 - Run `emctl stop dbconsole` to stop the Oracle Enterprise Manager agent.
 - Connect to the database instances with SQL*Plus as a privileged user and issue the `SHUTDOWN` command.
 - Connect to the Oracle ASM instance with SQL*Plus as a privileged user and issue the `SHUTDOWN` command.
 - Run `lsnrctl` and enter the `STOP` command to stop the listener.

For information about shutting down an Oracle ASM instance, see ["Shutting Down an Oracle ASM Instance"](#) on page 3-16.

See Also:

- Oracle Enterprise Manager manuals and online help for information about starting and stopping the Oracle Enterprise Manager agent
- *Oracle Database Administrator's Guide* for more information about starting up and shutting down Oracle instances
- *Oracle Database Net Services Administrator's Guide* for information about configuring a listener

2. Start the Oracle grid infrastructure OUI and select the **Upgrade Oracle grid infrastructure** option.
3. Complete the screens in the OUI installer and run the scripts as prompted by the OUI installer.

For example, on Linux you must run the `root.sh` script as the `root` user.

```
# GRID_HOME/root.sh
```

See Also: *Oracle Grid Infrastructure Installation Guide* for information about installing and upgrading Oracle grid infrastructure

4. Confirm that the listener and Oracle ASM instance are running in the Oracle grid infrastructure home and ensure that the Oracle Database instance and Oracle Enterprise Manager agent are running in the old database home.

- Confirm that the listener is running.

Otherwise start the listener with Server Control Utility (SRVCTL).

For example:

```
$ svrctl start listener
```

- Confirm that the Oracle ASM instance is running.

For example:

```
$ svrctl status listener
$ svrctl status asm
```

Otherwise start the Oracle ASM instance with SRVCTL.

For example:

```
$ svrctl start asm
```

- Ensure that the database instances are running; otherwise connect to the database instances with SQL*Plus as a privileged user and issue the `STARTUP` command.
- Ensure that the Oracle Enterprise Manager agent is running; otherwise start the Oracle Enterprise Manager agent with `emctl start dbconsole`.

For information about copying and moving an Oracle ASM instance initialization parameter file after upgrading, see ["Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File"](#) on page 3-4.

Note: The procedure described in this section upgrades the Oracle ASM instance only. Oracle Database, and Oracle Enterprise Manager, will not have the latest features. To upgrade Oracle Database, see *Oracle Database Upgrade Guide*.

Downgrading an Oracle ASM Instance

This section discusses the process to downgrade an Oracle ASM instance that has been upgraded to an Oracle Restart configuration.

Before you make any changes to the Oracle software, Oracle recommends that you create a backup of the Oracle software.

The following procedure describes how to downgrade an Oracle ASM instance from 11g release 2 (11.2) to 11g release 1 (11.1). In this scenario, the Oracle ASM instance was previously upgraded from an 11g release 1 (11.1) home to an Oracle Restart (single-instance) 11g release 2 (11.2) configuration. The 11g release 1 (11.1) home was not removed.

1. Determine disk group compatibility attribute settings.

If compatibility attributes have been advanced, then the disk groups must be recreated using compatibility attributes that allow access by the downgraded Oracle ASM and Oracle Database instances. A new disk group must be created with the old compatibility attributes and then you must restore the database files that were in the disk group.

When you revert to a new disk group with the old compatibility attribute settings, the latest Oracle ASM features might not be available. For example, if you revert the disk group compatibility to a pre-11.2 value, Oracle ACFS functionality is not available.

Copy or move an Oracle ASM SPFILE in a disk group to the file system before reverting disk group compatibility. Check the initialization parameters to ensure they are compatible with Oracle ASM 11g release 1 (11.1).

For information about reverting disk group compatibility, see ["Reverting Disk Group Compatibility"](#) on page 4-37. For information about moving data files between disk groups, see ["Moving Data Files Between Oracle ASM Disk Groups Using RMAN"](#) on page 8-9.

2. Downgrade any client databases from 11g release 2 (11.2) down to 11g release 1 (11.1).

See Also: *Oracle Database Upgrade Guide* for information about downgrading an Oracle Database and Oracle Enterprise Manager

3. Shut down the Oracle Enterprise Manager agent, Oracle Database instance, Oracle ASM instance, and the listener the in the database and Oracle ASM homes.

- Run `emctl stop dbconsole` to stop the Oracle Enterprise Manager agent.
- Connect to the database instances with SQL*Plus as a privileged user and issue the `SHUTDOWN` command.
- Shut down the Oracle ASM instance with Server Control Utility (SRVCTL).

```
$ svrctl stop asm
```

- Stop the listener with SVRCTL.

```
$ svrctl stop listener
```

For information about shutting down an Oracle ASM instance, see ["Shutting Down an Oracle ASM Instance"](#) on page 3-16.

See Also:

- Oracle Enterprise Manager manuals and online help for information about starting and stopping the Oracle Enterprise Manager agent
- *Oracle Database Administrator's Guide* for more information about starting up and shutting down Oracle instances
- *Oracle Database Net Services Administrator's Guide* for information about configuring a listener
- *Oracle Real Application Clusters Administration and Deployment Guide* for information about Server Control Utility (SRVCTL)

4. Deconfigure the Oracle Restart 11g release 2 (11.2) configuration.

Run the `roothas.pl` script as root.

For example, on Linux:

```
# GRID_HOME/crs/install/roothas.pl -delete
```

The 11g release 2 (11.2) `inittab` and `init*` scripts should be removed with the deconfiguration of Oracle Clusterware.

5. Unload the Oracle ACFS drivers.

For example, on Linux run `acfsload stop` as root.

```
# GRID_HOME/bin/acfsload stop
```

For information about Oracle ACFS driver resource management, see ["Oracle ACFS Drivers Resource Management"](#) on page B-3.

6. Recreate the Oracle ASM 11g release 1 (11.1) resources.

Run `localconfig` as root to add the resources to the Oracle ASM 11g release 1 (11.1) home.

For example, on Linux:

```
# ORACLE_ASM_11.1_HOME/bin/localconfig add
```

If `localconfig add` fails, use the `reset` option followed by the `ORACLE_HOME` to reset the existing resources.

For example, on Linux:

```
# localconfig reset ORACLE_ASM_11.1_HOME
```

7. Confirm that the Oracle ASM `PFILE` and `listener.ora` files are present in the Oracle ASM 11g release 1 (11.1) home.

If the Oracle ASM 11g release 1 (11.1) home has not been removed, the files should be available.

8. Configure additional configuration files in the Oracle ASM 11g release 1 (11.1) home.

For example, update files in the `/etc` directory on the Linux machine.

Update the Oracle ASM entry in `/etc/oratab` to point to the Oracle ASM 11g release 1 (11.1) home.

```
+ASM: /ORACLE_ASM_11.1.1_HOME/product/11.1.0/asm_1:N
```

The 11g release 2 (11.2) `inittab` and `init*` scripts should be removed with the deconfiguration of Oracle Clusterware.

9. Ensure that the listener, Oracle ASM instance, Oracle Database instance, and Oracle Enterprise Manager agent are running in the 11g release 1 (11.1) Oracle ASM and database homes.
 - Start the listener with `lsnrctl` and enter the `START` option.

If necessary, start Network Configuration Assistant (NETCA) in the Oracle ASM 11g release 1 (11.1) home with `netca`. Follow the prompts in the wizard to reconfigure the listener.
 - Connect to the Oracle ASM instance with SQL*Plus as a privileged user and issue the `STARTUP` command.
 - Connect to the database instances with SQL*Plus as a privileged user and issue the `STARTUP` command.
 - Start the Oracle Enterprise Manager agent with `emctl start dbconsole`.

Oracle Enterprise Manager may need to be reconfigured after the Oracle ASM instance has been downgraded.

See Also: *Oracle Grid Infrastructure Installation Guide* for information about installation and upgrading of Oracle grid infrastructure

Active Session History Sampling for Oracle ASM

Active Session History sampling is now available on Oracle ASM instances. This activity is exposed in the dynamic `V$ACTIVE_SESSION_HISTORY` view. Active Session History sampling requires a diagnostic pack license for the Oracle ASM instance.

See Also:

- *Oracle Database Performance Tuning Guide* for more information about gathering performance statistics
- *Oracle Database Reference* for a description of the `V$ACTIVE_SESSION_HISTORY` view

Using Oracle ASM Rolling Upgrade

Oracle ASM rolling upgrade enable you to independently upgrade or patch clustered Oracle ASM nodes without affecting database availability, thus providing greater uptime. Rolling upgrade means that some features of a clustered Oracle ASM environment continue to function when one or more of the nodes in the cluster uses different software versions.

Oracle recommends that you perform an Oracle ASM rolling upgrade when performing n Oracle Clusterware rolling upgrade.

To perform a rolling upgrade, your environment must be prepared. Oracle Clusterware must be fully upgraded to the next patch or release version before you start the Oracle ASM rolling upgrade. In addition, you should prepare your Oracle Clusterware in a rolling upgrade manner to ensure high availability and maximum

uptime. Note that the rolling upgrade to 11g Release 2 (11.2) moves the Oracle ASM instance to 11g Release 2 (11.2) Oracle grid infrastructure home.

You can upgrade a single Oracle ASM instance with Oracle Universal Installer (OUI). For information, see "[Upgrading an Oracle ASM Instance With Oracle Universal Installer](#)" on page 3-17.

Notes:

- Rolling upgrades only apply to clustered Oracle ASM instances, and you can only perform rolling upgrades on environments with Oracle Database 11g or later. In other words, you cannot use this feature to upgrade from Oracle Database 10g to Oracle Database 11g.
 - See *Oracle Exadata documentation* for information about performing a rolling upgrading of an Oracle ASM instance when Oracle Exadata storage is present.
-
-

See Also:

- *Oracle Grid Infrastructure Installation Guide* for information about performing a rolling upgrade of Oracle ASM
- *Oracle Database Upgrade Guide* for information upgrading an Oracle Database
- *Oracle Database SQL Language Reference* for information about the rolling migration clause of the `ALTER SYSTEM` commands

Patching Oracle ASM Instances

For Oracle RAC environments, ensure that your Oracle Clusterware version is at least equal to the version of the patch that you are applying to the Oracle Database. First apply the patch to the Oracle grid infrastructure home and then apply the patch to the Oracle Database home.

Note: You must apply the patch to the Oracle grid infrastructure home before you apply it to the Oracle Database home.

Authentication for Accessing Oracle ASM Instances

An Oracle ASM instance does not have a data dictionary, so the only way to connect to an Oracle ASM instance is by using one of three system privileges, `SYSASM`, `SYSDBA`, or `SYSOPER`. There are three modes of connecting to Oracle ASM instances:

- Local connection using operating system authentication
- Local connection using password authentication
- Remote connection by way of Oracle Net Services using password authentication

See Also: Your operating system-specific *Oracle Grid Infrastructure Installation Guide* for information about how to ensure that the Oracle ASM and database instances have member disk access

This section describes the following topics:

- [About Privileges for Oracle ASM](#)
- [Creating Users with the SYSASM Privilege](#)
- [Operating System Authentication for Oracle ASM](#)
- [Password File Authentication for Oracle ASM](#)

The Oracle ASM and database instances must have read/write operating system access rights to disk groups. For example, the Oracle ASM instance and the database instance must have identical read and write permissions for the disks that comprise the related Oracle ASM disk group. For Linux and UNIX systems, this is typically provided through shared Linux and UNIX group membership (OSASM group). On Windows systems, the Oracle ASM service must be run as Administrator. For information about file permissions and Oracle ASM File Access Control, see "[Managing Oracle ASM File Access Control for Disk Groups](#)" on page 4-38.

About Privileges for Oracle ASM

During Oracle ASM installation, you can use one operating system group for all users or divide system privileges so that database administrators, storage administrators, and database operators each have distinct operating system privilege groups.

Whether you create separate operating system privilege groups or use one group to provide operating system authentication for all system privileges, you should use SYSASM to administer an Oracle ASM instance. The SYSDBA privilege cannot be used to administer an Oracle ASM instance. If you use the SYSDBA privilege to run administrative commands on an Oracle ASM instance, the operation results in an error. The SYSDBA privilege is intended to be used by the database to access disk groups.

Oracle also recommends that the use of a less privileged user, such as ASMSNMP with SYSDBA privileges that is created during installation, for monitoring the Oracle ASM instance.

Operating system authentication using membership in the group or groups designated as OSDBA, OSOPER, and OSASM is valid on all Oracle platforms. Connecting to an Oracle ASM instance as SYSASM grants you full access to all of the available Oracle ASM disk groups and management functions.

This section contains these topics:

- [Using One Operating System Group for Oracle ASM Users](#)
- [Using Separate Operating System Groups for Oracle ASM Users](#)
- [The SYSASM Privilege for Administering Oracle ASM](#)
- [The SYSDBA Privilege for Managing Oracle ASM Components](#)

For information about privileges and Oracle ACFS, see "[Oracle ACFS and File Access and Administration Security](#)" on page 5-8.

Using One Operating System Group for Oracle ASM Users

If you do not want to divide system privileges access into separate operating system groups, then you can designate one operating system group as the group whose members are granted access as OSDBA, OSOPER, and OSASM for Oracle ASM privileges. The default operating system group name for all of these is usually dba and that group is typically chosen for the default configuration.

[Table 3–1](#) shows an example of a Linux deployment without separated privileges for Oracle ASM users.

Table 3–1 One Operating System Group and One Set of Privileges for All Oracle ASM Users

Role/Software Owner	User	Group/Privilege
Oracle ASM administrator/Oracle grid infrastructure home	oracle	dba/SYSASM, SYSDBA, SYSOPER
Database administrator 1/Database home 1	oracle	dba/SYSASM, SYSDBA, SYSOPER
Database administrator 2/Database home 2	oracle	dba/SYSASM, SYSDBA, SYSOPER
Operating system disk device owner	oracle	dba

Using Separate Operating System Groups for Oracle ASM Users

You can designate separate operating system groups as the operating system authentication groups for privileges on Oracle ASM.

The following list describes the separate operating system authentication groups for Oracle ASM and the privileges that their members are granted.

- OSASM group

This group is granted the SYSASM privilege, which provides full administrative privileges for the Oracle ASM instance. For example, the group could be `asmadmin`.
- OSDBA for Oracle ASM group

This group is granted the SYSDBA privilege on the Oracle ASM instance, which grants access to data stored on Oracle ASM. This group has a subset of the privileges of the OSASM group.

When you implement separate administrator privileges, choose an OSDBA group for the Oracle ASM instance that is different than the group that you select for the database instance, such as `dba`. For example, the group could be `asmdba`.
- OSOPER for Oracle ASM group

This group is granted the SYSOPER privilege on the Oracle ASM instance, which provides operations such as startup, shutdown, mount, dismount, and check disk group. This group has a subset of the privileges of the OSASM group. For example, the group could be `asmoper`.

When you implement separate Oracle ASM and database administrator duties, this configuration requires different group and different software owners. Implicitly this implementation requires that the OSASM and OSDBA are different groups. For this configuration, you must create an OSDBA for Oracle ASM group and a database instance must be a member of that group to access the Oracle ASM instance.

In an installation that has been configured as Oracle Clusterware, the Oracle ASM user, such as `grid`, does not need to be a member of the Oracle Database OSDBA group, such as `dba1` or `dba2`, because the Oracle Clusterware database agent runs as the database owner and can use SYSDBA to connect to the database.

However, in an Oracle Restart configuration, the Oracle ASM user (`grid`) needs to be a member of the OSDBA group (`dba1`, `dba2`, ...) of every database. This requirement is necessary because Oracle Restart software runs as the Oracle ASM user (`grid`) and this user must be able to start and stop the databases using the `CONNECT / AS SYSDBA` authentication.

Additionally, the owner of the operating system disk devices should be the same as owner of the Oracle ASM software.

Table 3–2 shows an example of a Linux deployment using separate operating system privilege groups for Oracle ASM users.

Table 3–2 Separated Operating System Groups and Privileges for Oracle ASM Users

Role/Software Owner	User	Group/Privilege
Oracle ASM administrator/Oracle grid infrastructure home	grid	asmadmin (OSASM)/SYSASM asmdba (OSDBA for ASM)/SYSDBA asmoper (OSOPER for ASM)/SYSOPER dba1, dba2, ... (OSDBA for the databases when in an Oracle Restart configuration)
Database administrator 1/Database home 1	oracle1	asmdba (OSDBA for ASM)/SYSDBA oper1 (OSOPER for database 1)/SYSOPER dba1 (OSDBA for database 1)/SYSDBA
Database administrator 2/Database home 2	oracle2	asmdba (OSDBA for ASM)/SYSDBA oper2 (OSOPER for database 2)/SYSOPER dba2 (OSDBA for database 2)/SYSDBA
Operating system disk device owner	grid	asmadmin (OSASM)

The SYSASM Privilege for Administering Oracle ASM

SYSASM is a system privilege that enables the separation of the SYSDBA database administration privilege from the Oracle ASM storage administration privilege. Access to the SYSASM privilege is granted by membership in an operating system group that is designated as the OSASM group. This is similar to SYSDBA and SYSOPER privileges, which are system privileges granted through membership in the groups designated as the OSDBA and OSOPER operating system groups. You can designate one group for all of these system privileges, or you can designate separate groups for each operating system privilege.

You can also grant the SYASM privilege with password file authentication, as discussed in "Password File Authentication for Oracle ASM" on page 3-27.

To connect locally as SYSASM using password authentication with SQL*Plus, use the following statement:

```
sqlplus SYS AS SYSASM
...
Enter password:
```

To connect remotely as SYSASM using password authentication with SQL*Plus, use the following statement:

```
sqlplus sys@\"myhost.mydomain.com:1521/+ASM\" AS SYSASM
...
Enter password:
```

In the previous example, +ASM is the service name of the Oracle ASM instance.

To connect locally as SYSASM to an Oracle ASM instance using operating system authentication with SQL*Plus, use the following statement:

```
sqlplus / AS SYSASM
```

The SYSDBA Privilege for Managing Oracle ASM Components

You can connect as SYSDBA to use SQL*Plus or ASMCMD commands to manage Oracle ASM components associated with the database. When running SQL or ASMCMD operations with the SYSDBA privilege, connect to the database instance rather than the Oracle ASM instance.

Connecting as SYSDBA to the database instance has limited set of Oracle ASM privileges. For example, you cannot create a disk group when connected with the SYSDBA privilege.

When connected as SYSDBA to the database instance, the Oracle ASM operations are limited to:

- Create and delete files, aliases, directories, and templates
- Examine various Oracle ASM instance views
- Operate on files that were created by this user or only access files to which another user had explicitly granted access
- Granting Oracle ASM File Access Control to other users

Creating Users with the SYSASM Privilege

When you are logged in to an Oracle ASM instance as SYSASM, you can use the combination of CREATE USER and GRANT SQL statements to create a new user who has the SYSASM privilege. You also can revoke the SYSASM privilege from a user using the REVOKE command, and you can drop a user from the password file using the DROP USER command.

Note: These commands update the password file for the local Oracle ASM instance only.

The following example describes how to perform these SQL operations for the user identified as *new_user*:

```
REM create a new user, then grant the SYSASM privilege
SQL> CREATE USER new_user IDENTIFIED by new_user_passwd;
SQL> GRANT SYSASM TO new_user;
```

```
REM connect the user to the ASM instance
SQL> CONNECT new_user AS SYSASM;
Enter password:
```

```
REM revoke the SYSASM privilege, then drop the user
SQL> REVOKE SYSASM FROM new_user;
SQL> DROP USER new_user;
```

For information about creating a user with Oracle ASM command-line utility (ASMCMD), see ["orapwusr"](#) on page 12-13. For information about creating a user with Oracle Enterprise Manager, see ["Managing Oracle ASM Users with Oracle Enterprise Manager"](#) on page 9-6.

Operating System Authentication for Oracle ASM

Membership in the operating system group designated as the OSASM group provides operating system authentication for the SYSASM system privilege. OSASM is provided exclusively for Oracle ASM. Initially, only the user that installs ASM is a

member of the OSASM group, if you use a separate operating system group for that privilege. However, you can add other users. Members of the OSASM group are authorized to connect using the SYSASM privilege and have full access to Oracle ASM, including administrative access to all disk groups that are managed by that Oracle ASM instance.

On Linux and UNIX systems, the default operating system group designated as OSASM, OSOPER for Oracle ASM, and OSDBA for Oracle ASM is `dba`. On Windows systems, the default name designated as OSASM, OSOPER, and OSDBA is `ora_dba`.

SQL*Plus commands, ASMCMD commands, and ASMCA use operating system authentication.

See Also:

- *Oracle Database Administrator's Guide* for more information about using operating system authentication
- *Oracle Grid Infrastructure Installation Guide* for information about installation of the Oracle grid infrastructure

Password File Authentication for Oracle ASM

Password file authentication for Oracle ASM can work both locally and remotely. To enable password file authentication, you must create a password file for Oracle ASM. A password file is also required to enable Oracle Enterprise Manager to connect to Oracle ASM remotely.

If you select the Oracle ASM storage option, then ASMCA creates a password file for Oracle ASM when it initially configures the Oracle ASM disk groups. Similar to a database password file, the only user added to the password file when ASMCA creates it is SYS. To add other users to the password file, you can use the CREATE USER and GRANT commands as described previously in the section titled "[About Privileges for Oracle ASM](#)" on page 3-23.

If you configure an Oracle ASM instance without using ASMCA, then you must manually create a password file and grant the SYSASM privilege to user SYS.

SQL*Plus commands and Oracle Enterprise Manager use password file authentication.

See Also:

- *Oracle Database Administrator's Guide* for information about creating and maintaining a password file
- *Oracle Database SQL Language Reference* for information about the CREATE USER and GRANT commands
- *Oracle Database Security Guide* for information about database security
- *Oracle Database Reference* for information about the V\$PWFILE_USERS view which lists users who have been granted SYSASM, SYSDBA, and SYSOPER privileges as derived from the password file.

Migrating a Database to Use Oracle ASM

With a new installation of Oracle Database and Oracle ASM, you can initially create your database and select the Oracle ASM storage option. If you have an existing

Oracle database that stores database files in the operating system file system or on raw devices, then you can migrate some or all of your data files to Oracle ASM storage.

Oracle provides several methods for migrating your database to Oracle ASM. Using Oracle ASM enables you to realize the benefits of automation and simplicity in managing your database storage. You can use the following methods to migrate to Oracle ASM as described in this section:

- [Using Oracle Enterprise Manager to Migrate Databases to Oracle ASM](#)
- [Using Oracle Recovery Manager to Migrate Databases to Oracle ASM](#)
- [Best Practices White Papers on Migrating to Oracle ASM](#)

Note: You must upgrade to at least Oracle Database 10g before migrating your database to Oracle ASM.

Using Oracle Enterprise Manager to Migrate Databases to Oracle ASM

Oracle Enterprise Manager enables you to perform cold and hot database migration with a GUI. You can access the migration wizard from the Oracle Enterprise Manager Home page under the Change Database heading.

For more information about using Oracle Enterprise Manager to upgrade to Oracle ASM, see [Chapter 9, "Administering Oracle ASM with Oracle Enterprise Manager"](#).

Using Oracle Recovery Manager to Migrate Databases to Oracle ASM

You can use Oracle Recovery Manager (RMAN) to manually migrate to Oracle ASM. You can also use RMAN to migrate a single tablespace or data file to Oracle ASM.

For more information, see [Chapter 8, "Performing Oracle ASM Data Migration With RMAN"](#).

Best Practices White Papers on Migrating to Oracle ASM

The Oracle Maximum Availability Architecture (MAA) Web site provides excellent best practices technical white papers based on different scenarios, such as:

- Minimal Downtime Migration to Oracle ASM
- Platform Migration using Transportable Tablespaces
- Platform Migration using Transportable Database

See Also: For information about Oracle ASM best practices for migrating to Oracle ASM from environments that do not use Oracle ASM, refer to the following MAA link on OTN:

<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>

Administering Oracle ASM Disk Groups

This chapter describes how to administer Oracle Automatic Storage Management (Oracle ASM) disk groups. This information includes how to create, alter, drop, mount, and dismount Oracle ASM disk groups. The database instances that use Oracle ASM can continue operating while you administer disk groups.

The examples in this chapter use SQL statements. These examples assume that SQL*Plus is run from the Oracle grid home where Oracle ASM is installed and the Oracle environmental variables are set to this home. The examples also assume that the Oracle ASM instance is running. This chapter contains the following topics:

- [Disk Group Attributes](#)
- [Creating Disk Groups](#)
- [Altering Disk Groups](#)
- [Oracle ASM Disk Discovery](#)
- [Managing Capacity in Disk Groups](#)
- [Oracle ASM Mirroring and Disk Group Redundancy](#)
- [Performance and Scalability Considerations for Disk Groups](#)
- [Disk Group Compatibility](#)
- [Managing Oracle ASM File Access Control for Disk Groups](#)
- [Mounting and Dismounting Disk Groups](#)
- [Checking the Internal Consistency of Disk Group Metadata](#)
- [Dropping Disk Groups](#)
- [Renaming Disks Groups](#)

For information about starting up an Oracle ASM instance, refer to ["Starting Up an Oracle ASM Instance"](#) on page 3-13.

For information about administering Oracle ASM disk groups with Oracle Enterprise Manager, refer to [Chapter 9, "Administering Oracle ASM with Oracle Enterprise Manager"](#).

For information about administering Oracle ASM disk groups with Oracle ASM Configuration Assistant (ASMCA), refer to [Chapter 11, "Oracle ASM Configuration Assistant"](#).

For information about administering Oracle ASM disk groups with ASMCMD, refer to [Chapter 12, "Oracle ASM Command-Line Utility"](#).

See Also: The Oracle ASM home page for more information about Oracle ASM best practices at:
<http://www.oracle.com/technology/products/database/asm/index.html>

Disk Group Attributes

Disk group attributes are parameters that are bound to a disk group, rather than an Oracle ASM instance.

Disk group attributes can be set when a disk group is created or altered, unless otherwise noted in the following list.

- `ACCESS_CONTROL.ENABLED`

This attribute can only be set when altering a disk group.

For information about the `ACCESS_CONTROL.ENABLED` attribute, see "Using SQL Statements to Set Disk Group Attributes for Oracle ASM File Access Control" on page 4-40.

- `ACCESS_CONTROL.UMASK`

This attribute can only be set when altering a disk group.

For information about the `ACCESS_CONTROL.UMASK` attribute, see "Using SQL Statements to Set Disk Group Attributes for Oracle ASM File Access Control" on page 4-40.

- `AU_SIZE`

This attribute can only be set when creating a disk group.

For information about allocation unit size and extents, see "Allocation Units" on page 1-7, "Extents" on page 1-8, and "Specifying the Allocation Unit Size" on page 4-7. For an example of the use of the `AU_SIZE` attribute, see Example 4-1, "Creating the DATA Disk Group" on page 4-5.

- `CELL.SMART_SCAN_CAPABLE`

This attribute is only applicable to Oracle Exadata storage.

See Also: *Oracle Exadata documentation* for information about the `CELL.SMART_SCAN_CAPABLE` attribute

- `COMPATIBLE.ASM`

For information about the `COMPATIBLE.ASM` attribute, see "COMPATIBLE.ASM" on page 4-33.

- `COMPATIBLE.RDBMS`

For information about the `COMPATIBLE.RDBMS` attribute, see "COMPATIBLE.RDBMS" on page 4-33.

- `COMPATIBLE.ADVM`

For information about the `COMPATIBLE.ADVM` attribute, see "COMPATIBLE.ADVM" on page 4-33 and "Overview of Oracle ASM Dynamic Volume Manager" on page 5-10.

- `DISK_REPAIR_TIME`

This attribute can only be set when altering a disk group.

For information about the `DISK_REPAIR_TIME` attribute, see ["Oracle ASM Fast Mirror Resync"](#) on page 4-27.

- `SECTOR_SIZE`

This attribute can only be set when creating a disk group.

For information about the `SECTOR_SIZE` attribute, see ["Specifying the Sector Size for Drives"](#) on page 4-8.

In addition to the disk group attributes listed in this section, template attributes are also assigned to a disk group. For information about template attributes, see ["Managing Disk Group Templates"](#) on page 7-15.

You can display disk group attributes with the `V$ASM_ATTRIBUTE` view and the `ASMCMD lsattr` command. For an example of the use of the `V$ASM_ATTRIBUTE` view, see [Example 6-1](#) on page 6-2. For information about the `lsattr` command, see ["lsattr"](#) on page 12-33.

Creating Disk Groups

This section contains information about creating disk groups. You can use the `CREATE DISKGROUP` SQL statement to create a disk group.

This section contains the following topics:

- [Using the CREATE DISKGROUP SQL Statement](#)
- [Example: Creating a Disk Group](#)
- [Creating Disk Groups for a New Oracle Installation](#)
- [Specifying the Allocation Unit Size](#)
- [Specifying the Sector Size for Drives](#)
- [Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#)

Using the CREATE DISKGROUP SQL Statement

The `CREATE DISKGROUP` SQL statement is used to create disk groups. When creating a disk group, you:

- Assign a unique name to the disk group.

The specified disk group name is not case sensitive and is always converted to uppercase when stored internally.

Note: Oracle does not recommend using identifiers for database object names that must be quoted. While these quoted identifiers may be valid as names in the SQL `CREATE` statement, such as `CREATE DISKGROUP "1DATA"`, the names may not be valid when using other tools that manage the database object.

- Specify the redundancy level of the disk group.

If you want Oracle ASM to mirror files, specify the redundancy level as `NORMAL REDUNDANCY` (2-way mirroring by default for most file types) or `HIGH REDUNDANCY` (3-way mirroring for all files). You specify `EXTERNAL REDUNDANCY` if you do not want mirroring by Oracle ASM. For example, you might choose `EXTERNAL REDUNDANCY` if you want to use storage array protection features.

After a disk group is created, you cannot alter the redundancy level of the disk group. To change the redundancy level, you must create another disk group with the appropriate redundancy and then move the files to the new disk group.

Oracle recommends that you create failure groups of equal size to maintain space balance and even distribution of mirror data.

For more information about redundancy levels, refer to "[Oracle ASM Mirroring and Failure Groups](#)" on page 4-23.

- Specify the disks that are to be formatted as Oracle ASM disks belonging to the disk group.

The disks can be specified using operating system dependent wildcard characters in search strings that Oracle ASM then uses to find the disks. You can specify names for the disks with the `NAME` clause or use the system-generated names.

- Optionally specify the disks as belonging to specific failure groups.

For information about failure groups, refer to "[Understanding Oracle ASM Concepts](#)" on page 1-2 and "[Oracle ASM Mirroring and Failure Groups](#)" on page 4-23.

- Optionally specify the type of failure group.

For information about `QUORUM` and `REGULAR` failure groups, refer to "[Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#)" on page 4-10.

- Optionally specify disk group attributes, such software compatibility or allocation unit size.

Oracle ASM programmatically determines the size of each disk. If for some reason this is not possible, or if you want to restrict the amount of space used on a disk, you are able to specify a `SIZE` clause for each disk. Oracle ASM creates operating system-independent names for the disks in a disk group that you can use to reference the disks in other SQL statements. Optionally, you can provide your own name for a disk using the `NAME` clause. Disk names are available in the `V$ASM_DISK` view.

Note: A disk cannot belong to multiple disk groups.

The Oracle ASM instance ensures that any disk in a newly created disk group is addressable and is not currently a member of another disk group. You must use `FORCE` only when adding a disk that was dropped `FORCE`. If a disk is dropped `NOFORCE`, then use can add it `NOFORCE`. For example, a disk might have failed and was dropped from its disk group. After the disk is repaired, it is no longer part of any disk group, but Oracle ASM still recognizes that the disk had been a member of a disk group. You must use the `FORCE` flag to include the disk in a new disk group. In addition, the disk must be addressable, and the original disk group must not be mounted. Otherwise, the operation fails.

Note: Use caution when using the `FORCE` option to add a previously used disk to a disk group; you might cause another disk group to become unusable.

The `CREATE DISKGROUP` statement mounts the disk group for the first time, and adds the disk group name to the `ASM_DISKGROUPS` initialization parameter if a server parameter file is being used. If a text initialization parameter file is being used and you want the disk group to be automatically mounted at instance startup, then you must

remember to add the disk group name to the `ASM_DISKGROUPS` initialization parameter before the next time that you shut down and restart the Oracle ASM instance. You can also create disk groups with Oracle Enterprise Manager. Refer to ["Creating Disk Groups"](#) on page 9-7.

See Also: The `CREATE DISKGROUP` SQL statement in the *Oracle Database SQL Language Reference*

Example: Creating a Disk Group

The following examples assume that the `ASM_DISKSTRING` initialization parameter is set to the `'/devices/*'` string. Oracle ASM disk discovery identifies disks in the `/devices` directory, including the following disks:

Controller 1:

```
/devices/diska1
/devices/diska2
/devices/diska3
/devices/diska4
```

Controller 2:

```
/devices/diskb1
/devices/diskb2
/devices/diskb3
/devices/diskb4
```

The SQL statement in [Example 4-1](#) creates a disk group named `data` with normal redundancy consisting of two failure groups `controller1` or `controller2` with four disks in each failure group. The `data` disk group is typically used to store database data files.

Example 4-1 Creating the DATA Disk Group

```
CREATE DISKGROUP data NORMAL REDUNDANCY
  FAILGROUP controller1 DISK
    '/devices/diska1' NAME diska1,
    '/devices/diska2' NAME diska2,
    '/devices/diska3' NAME diska3,
    '/devices/diska4' NAME diska4
  FAILGROUP controller2 DISK
    '/devices/diskb1' NAME diskb1,
    '/devices/diskb2' NAME diskb2,
    '/devices/diskb3' NAME diskb3,
    '/devices/diskb4' NAME diskb4
  ATTRIBUTE 'au_size'='4M',
    'compatible.asm' = '11.2',
    'compatible.rdbms' = '11.2',
    'compatible.advm' = '11.2';
```

In [Example 4-1](#), the `NAME` clauses enable you to explicitly assign names to the disks rather than the default system-generated names. The system-generated names are in the form `diskgroup_nnnn`, where `nnnn` is the disk number for the disk in the disk group. For ASMLIB disks, the disk name defaults to the ASMLIB name that is the user label of the disk; for example, `mydisk` is the default Oracle ASM disk name for `ORCL:mydisk`.

When creating the disk group in [Example 4-1](#), the values of following disk group attributes were explicitly set:

- `AU_SIZE`
Specifies the size of the allocation unit for the disk group. For information about allocation unit size and extents, see "Extents" on page 1-8.
You can view the value of the `AU_SIZE` disk group attribute in the `ALLOCATION_UNIT_SIZE` column of the `V$ASM_DISKGROUP` view.
- `COMPATIBLE.ASM`
Determines the minimum software version for any Oracle ASM instance that uses a disk group. For information about the `COMPATIBLE.ASM` attribute, see "COMPATIBLE.ASM" on page 4-33.
- `COMPATIBLE.RDBMS`
Determines the minimum software version for any database instance that uses a disk group. For information about the `COMPATIBLE.RDBMS` attribute, see "COMPATIBLE.RDBMS" on page 4-33.
- `COMPATIBLE.ADVM`
Determines whether the disk group can contain Oracle ASM volumes. For information about the `COMPATIBLE.ADVM` attribute, see "COMPATIBLE.ADVM" on page 4-33.

In [Example 4-2](#), the `fra` disk group (typically created for the fast recovery area) is created with the default disk group attribute values. Names are not specified for the Oracle ASM disks and failure groups are not explicitly specified. This examples assumes that `diskc1` through `diskc9` are present in the `/devices` directory.

Example 4-2 Creating the FRA Disk Group

```
CREATE DISKGROUP fra NORMAL REDUNDANCY
  DISK '/devices/diskc*';
```

See Also: For information about using ASMLIB when creating disk groups, refer to the Oracle ASMLib page on the Oracle Technology Network Web site at <http://www.oracle.com/technology/tech/linux/asmlib/index.html>

Creating Disk Groups for a New Oracle Installation

This section describes the basic procedure to follow when creating disk groups during a new installation of Oracle Restart and Oracle Database. This information also applies to an Oracle grid infrastructure installation.

The procedure assumes that the `data` disk group is used for the storage of the database data files and the `fra` disk group is used for storage of the fast recovery area files. Detailed information about installation with Oracle Universal Installer (OUI) and database creation with Database Configuration Assistant (DBCA) is available in the installation guides for your specific operating system.

1. Install Oracle Restart with OUI, following the screen prompts.

During the installation, create the `data` disk group for storing database files such as the data and control files.

This OUI process is similar to creating a disk group with Oracle ASM Configuration Assistant (ASMCA). For information about using ASMCA to create

disk groups, see ["Managing Disk Groups With Oracle ASM Configuration Assistant"](#) on page 11-6.

Note that the `data` disk group is the disk group used for storing Oracle Cluster Registry (OCR) and voting disks in an Oracle grid infrastructure installation. See ["Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups"](#) on page 4-10

2. After Oracle Restart is installed, use ASMCA to create the `fra` disk group for storing the fast recovery area files.

Create the `fra` disk group to hold the fast recovery area files.

At this time, you can also update the `data` disk group if necessary. For information about using ASMCA to create or alter disk groups, see ["Managing Disk Groups With Oracle ASM Configuration Assistant"](#) on page 11-6.

You can also create the `fra` disk group with SQL*Plus or ASMCMD commands run from the Oracle Restart home. For information, see ["Using the CREATE DISKGROUP SQL Statement"](#) on page 4-3 and ["mkdg"](#) on page 12-40.

See Also:

- *Oracle Database Backup and Recovery User's Guide* for information about configuring the fast recovery area
- *Oracle Database Administrator's Guide* for information about specifying a fast recovery area

3. Install the Oracle Database software with OUI, following the screen prompts.
4. After the database software has been installed, run DBCA to create a database, following the screen prompts.

During the creation of the database, make the appropriate selections to use Oracle ASM for storage of data files and fast recovery area files. When prompted:

- Store database data files in the `data` disk group
- Store fast recovery area files in the `fra` disk group

See Also:

- *Oracle grid infrastructure Installation Guide* for your operating system for detailed information about installation of Oracle grid infrastructure and Oracle Restart
- *Oracle Database Installation Guide* for your operating system for detailed information about installation of an Oracle Database

Specifying the Allocation Unit Size

Oracle recommends that the allocation unit (AU) size for a disk group be set to 4 megabytes (MB). In addition to this AU size recommendation, the operating system (OS) I/O size should be set to the largest possible size.

Some of the benefits with a 4 MB allocation unit are:

- Increased I/O through the I/O subsystem if the I/O size is increased to the AU size.
- Reduced SGA size to manage the extent maps in the database instance.
- Faster datafile initialization if the I/O size is increased to the AU size.

- Increased file size limits.
- Reduced database open time.

The allocation unit size is specified with the disk group attribute `AU_SIZE`. The AU size cannot be altered after a disk group is created. [Example 4-1](#) on page 4-5 shows how the `AU_SIZE` is specified with the `CREATE DISKGROUP` SQL statement.

For applications that can benefit from large I/O sizes, you can increase the `DB_FILE_MULTIBLOCK_READ_COUNT` initialization parameter to 4 MB from the default of 1 MB. Ensure that the application is thoroughly tested with this new setting as this change could adversely affect stable optimized plans and future optimizer plan generation.

To change the `DB_FILE_MULTIBLOCK_READ_COUNT` setting, use the `ALTER SYSTEM` SQL statement, as shown in [Example 4-3](#). This example assumes an 8 K block size.

Example 4-3 Altering the `DB_FILE_MULTIBLOCK_READ_COUNT` Setting

```
SQL> ALTER SYSTEM SET db_file_multiblock_read_count=512 scope=both sid=*
```

See Also: *Oracle Database Reference* for more information about the `DB_FILE_MULTIBLOCK_READ_COUNT` initialization parameter

Specifying the Sector Size for Drives

The optional `SECTOR_SIZE` disk group attribute can be used with the `CREATE DISKGROUP` SQL statement to specify disks with the sector size set to the value of `SECTOR_SIZE` for the disk group. Oracle ASM provides support for 4 KB sector disk drives without negatively affecting performance. The `SECTOR_SIZE` disk group attribute can be set only during disk group creation.

The values for `SECTOR_SIZE` can be 512, 4096, or 4K. The default value is platform dependent. The `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` disk group attributes must be set to 11.2 or higher to set the sector size to a value other than the default value.

Note: Oracle Automatic Storage Management Cluster File System (Oracle ACFS) does not support 4 KB sector drives. There is a performance penalty for Oracle ACFS when using 4 KB sector disk drives in 512 sector emulation mode.

The following validations apply to the sector size disk group attribute.

- Oracle ASM prevents disks of different sector sizes from being added to the same disk group. This validation occurs during `CREATE DISKGROUP`, `ALTER DISKGROUP ADD DISK`, and `ALTER DISKGROUP MOUNT` operations.
- If the `SECTOR_SIZE` attribute is explicitly specified when creating a disk group, then Oracle ASM attempts to verify that all disks discovered through disk search strings have a sector size equal to the specified value. If one or more disks were found to have sector size different from the specified value, or if Oracle ASM was not able to verify a disk sector size, then the create operation fails.

Oracle ASM also attempts to verify disk sector size during the mount operation and the operation fails if one or more disks have a sector size different than the value of the `SECTOR_SIZE` attribute.

- If the `SECTOR_SIZE` attribute is not specified when creating a disk group and Oracle ASM is able to verify that all discovered disks have the same sector value, then that value is assumed for the disk group sector size that is created. If the disks have different sector sizes, the create operation fails.
- When new disks are added to an existing disk group using the `ALTER DISKGROUP .. ADD DISK SQL` statement, you must ensure that the new disks that you want to add have the same value as the `SECTOR_SIZE` disk group attribute. If the new disks have different sector sizes, the alter operation fails.
- You can determine the sector size value that has either been assumed or explicitly set for a successful disk group creation by querying the `V$ASM_ATTRIBUTE` view or run the `ASMCMD lsattr` command. You can also query the `SECTOR_SIZE` column in the `V$ASM_DISKGROUP` view.

```
SQL> SELECT name, value FROM V$ASM_ATTRIBUTE
       WHERE name = 'sector_size' AND group_number = 1;
NAME                                VALUE
-----
sector_size                          512

SQL> SELECT group_number, sector_size FROM V$ASM_DISKGROUP
       WHERE group_number = 1;

GROUP_NUMBER SECTOR_SIZE
-----
1             512
```

As shown in [Example 4-4](#), you can use the `SECTOR_SIZE` attribute with the `CREATE DISKGROUP SQL` statement to specify the sector size of the disk drive on which the Oracle ASM disk group is located.

Example 4-4 Creating a Disk Group of 4K Sector Size

```
CREATE DISKGROUP data NORMAL REDUNDANCY
FAILGROUP controller1 DISK
'/devices/diska1',
'/devices/diska2',
'/devices/diska3',
'/devices/diska4'
FAILGROUP controller2 DISK
'/devices/diskb1',
'/devices/diskb2',
'/devices/diskb3',
'/devices/diskb4'
ATTRIBUTE 'compatible.asm' = '11.2', 'compatible.rdbms' = '11.2',
'sector_size'='4096';
```

See Also:

- *Oracle Database SQL Language Reference* for information about the disk group attributes and the `CREATE DISKGROUP SQL` statement
- *Oracle Database SQL Language Reference* for information about the `BLOCKSIZE` keyword
- *Oracle Database Reference* for details about the `V$ASM_ATTRIBUTE` and `V$ASM_DISKGROUP` views

Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups

You can store Oracle Cluster Registry (OCR) and the voting file in Oracle ASM disk groups. The voting file and OCR are two important components of Oracle Clusterware.

The voting file helps you to manage information about node membership. OCR is a system that manages cluster and Oracle Real Application Clusters (Oracle RAC) database configuration information. A quorum failure group is a special type of failure group and disks in these failure groups do not contain user data and are not considered when determining redundancy requirements. For information about failure groups, see "[Oracle ASM Failure Groups](#)" on page 4-24.

You can manage and monitor this feature with the following:

- The CRSCTL and ocrconfig command-line tools
The CRSCTL and ocrconfig commands enable the placement of OCR storage and Cluster Synchronization Services (CSS) voting disks inside the disk groups managed by Oracle ASM.
- CREATE/ALTER DISKGROUP SQL Statements
The SQL keywords QUORUM and REGULAR enable the specification of disk and failure groups when creating or altering disk groups.
See [Example 4-5](#) on page 4-11.
- V\$ASM views
The FAILGROUP_TYPE column in both the V\$ASM_DISK and V\$ASM_DISK_STAT views specifies failure group type. The value for this column is REGULAR for regular failure groups and QUORUM for quorum failure groups.
The VOTING_FILE column in both the V\$ASM_DISK and V\$ASM_DISK_STAT views specifies whether a disk contains a voting disk. The value for this column is Y if the disk contains a voting disk or N otherwise.
Note that the value of USABLE_FILE_MB in V\$ASM_DISKGROUP and V\$ASM_DISKGROUP_STAT does not consider any free space that is present in QUORUM disks because that space is not available for client data files.
See "[Views Containing Oracle ASM Disk Group Information](#)" on page 6-1.

The QUORUM and REGULAR keywords provide an additional qualifier for failure group or disk specifications when creating or altering a disk group. QUORUM disks (or disks in QUORUM failure groups) cannot have client data files, whereas REGULAR disks (or disks in non-quorum failure groups) have no such restriction.

These keywords can be used before the keyword FAILGROUP if a failure group is being explicitly specified. If the failure group is implicitly implied, these keywords (QUORUM/REGULAR) can be used before the keyword DISK. When failure groups are explicitly specified, it is an error to specify these keywords (QUORUM/REGULAR) immediately before the keyword DISK. REGULAR is the default failure group type.

When performing operations on existing disks or failure groups, the qualifier specified in the SQL must match the qualifier that was specified when the disks or failure groups were added to the disk group.

[Example 4-5](#) shows the creation of a disk group with a QUORUM failure group. For Oracle Clusterware files a minimum of three disk devices or three failure groups are required with a normal redundancy disk group. Also, a QUORUM failure group does not count towards redundancy requirements.

The `COMPATIBLE.ASM` disk group compatibility attribute must be set to 11.2 or greater to store OCR or voting disk data in a disk group.

Example 4–5 Using the QUORUM Keyword

```
CREATE DISKGROUP ocr_data NORMAL REDUNDANCY
  FAILGROUP fg1 DISK '/devices/diskg1'
  FAILGROUP fg2 DISK '/devices/diskg2'
  QUORUM FAILGROUP fg3 DISK '/devices/diskg3'
  ATTRIBUTE 'compatible.asm' = '11.2.0.0.0';
```

See Also:

- *Oracle Grid Infrastructure Installation Guide* for information about configuring voting disks and Oracle Cluster Registry (OCR)
- *Oracle Clusterware Administration and Deployment Guide* for information about administering voting disks and Oracle Cluster Registry (OCR) using `CRSCTL` and `ocrconfig` commands
- *Oracle Database SQL Language Reference* for information about the `CREATE DISKGROUP` and `ALTER DISKGROUP` SQL statements
- *Oracle Database Reference* for details about the `V$ASM*` dynamic performance views

Altering Disk Groups

You can use the `ALTER DISKGROUP` SQL statement to alter a disk group configuration. You can add, resize, or drop disks while the database remains online. Whenever possible, multiple operations in a single `ALTER DISKGROUP` statement are recommended.

Oracle ASM automatically rebalances when the configuration of a disk group changes. By default, the `ALTER DISKGROUP` statement does not wait until the operation is complete before returning. Query the `V$ASM_OPERATION` view to monitor the status of this operation.

You can use the `REBALANCE WAIT` clause if you want the `ALTER DISKGROUP` statement processing to wait until the rebalance operation is complete before returning. This is especially useful in scripts. The statement also accepts a `REBALANCE NOWAIT` clause that invokes the default behavior of conducting the rebalance operation asynchronously in the background.

You can interrupt a rebalance running in wait mode by typing `CTRL-C` on most platforms. This causes the statement to return immediately with the message `ORA-01013: user requested cancel of current operation`, and then to continue the operation asynchronously. Typing `CTRL-C` does not cancel the rebalance operation or any disk add, drop, or resize operations.

To control the speed and resource consumption of the rebalance operation, you can include the `REBALANCE POWER` clause in statements that add, drop, or resize disks. Refer to "[Manually Rebalancing Disk Groups](#)" on page 4-17 for more information about this clause.

This section contains the following topics:

- [Adding Disks to a Disk Group](#)
- [Adding Volumes to Disk Groups](#)

- [Dropping Disks from Disk Groups](#)
- [Intelligent Data Placement](#)
- [Resizing Disks in Disk Groups](#)
- [Undropping Disks in Disk Groups](#)
- [Manually Rebalancing Disk Groups](#)
- [Tuning Rebalance Operations](#)

See Also: The `ALTER DISKGROUP` SQL statement in the *Oracle Database SQL Language Reference*

Adding Disks to a Disk Group

You can use the `ADD` clause of the `ALTER DISKGROUP` statement to add a disk or a failure group to a disk group. The same syntax that you use to add a disk or failure group with the `CREATE DISKGROUP` statement can be used with the `ALTER DISKGROUP` statement. For an example of the `CREATE DISKGROUP` SQL statement, refer to [Example 4-1](#) on page 4-5. After you add new disks, the new disks gradually begin to accommodate their share of the workload as rebalancing progresses.

Oracle ASM behavior when adding disks to a disk group is best illustrated through "[Example: Adding Disks to a Disk Group](#)" on page 4-12. You can also add disks to a disk group with Oracle Enterprise Manager, described in "[Adding Disks to Disk Groups](#)" on page 9-9.

Example: Adding Disks to a Disk Group

The statements presented in this example demonstrate the interactions of disk discovery with the `ADD DISK` operation.

Assume that disk discovery identifies the following disks in directory `/devices`:

```
/devices/diska1 -- member of data1
/devices/diska2 -- member of data1
/devices/diska3 -- member of data1
/devices/diska4 -- member of data1
/devices/diska5 -- candidate disk
/devices/diska6 -- candidate disk
/devices/diska7 -- candidate disk
/devices/diska8 -- candidate disk

/devices/diskb1 -- member of data1
/devices/diskb2 -- member of data1
/devices/diskb3 -- member of data1
/devices/diskb4 -- member of data2

/devices/diskc1 -- member of data2
/devices/diskc2 -- member of data2
/devices/diskc3 -- member of data3
/devices/diskc4 -- candidate disk

/devices/diskd1 -- candidate disk
/devices/diskd2 -- candidate disk
/devices/diskd3 -- candidate disk
/devices/diskd4 -- candidate disk
/devices/diskd5 -- candidate disk
```



```

/devices/diskd6 -- candidate disk
/devices/diskd7 -- candidate disk
/devices/diskd8 -- candidate disk

```

You can query the `V$ASM_DISK` view to display the status of Oracle ASM disks. See ["Views Containing Oracle ASM Disk Group Information"](#) on page 6-1.

The following statement would fail because `/devices/diska1 - /devices/diska4` already belong to `data1`.

```

ALTER DISKGROUP data1 ADD DISK
    '/devices/diska*';

```

The following statement would successfully add disks `/devices/diska5` through `/devices/diska8` to `data1`. Because no `FAILGROUP` clauses are included in the `ALTER DISKGROUP` statement, each disk is assigned to its own failure group. The `NAME` clauses assign names to the disks, otherwise they would have been assigned system-generated names.

```

ALTER DISKGROUP data1 ADD DISK
    '/devices/diska5' NAME diska5,
    '/devices/diska6' NAME diska6,
    '/devices/diska7' NAME diska7,
    '/devices/diska8' NAME diska8;

```

The following statement would fail because the search string matches disks that are contained in other disk groups. Specifically, `/devices/diska4` belongs to disk group `data1` and `/devices/diskb4` belongs to disk group `data2`.

```

ALTER DISKGROUP data1 ADD DISK
    '/devices/disk*4';

```

The following statement would successfully add `/devices/diskd1` through `/devices/diskd8` to disk group `data1`. This statement runs with a rebalance power of 5, and does not return until the rebalance operation is complete.

```

ALTER DISKGROUP data1 ADD DISK
    '/devices/diskd*'
    REBALANCE POWER 5 WAIT;

```

If `/devices/diskc3` was previously a member of a disk group that no longer exists, then you could use the `FORCE` option to add them as members of another disk group. For example, the following use of the `FORCE` clause enables `/devices/diskc3` to be added to `data2`, even though it is a current member of `data3`. For this statement to succeed, `data3` cannot be mounted.

```

ALTER DISKGROUP data2 ADD DISK
    '/devices/diskc3' FORCE;

```

Adding Volumes to Disk Groups

You can create an Oracle ASM Dynamic Volume Manager (Oracle ADVM) volume in a disk group. The volume device associated with the dynamic volume can then be used to host an Oracle ACFS file system.

The compatibility parameters `COMPATIBLE.ASM` and `COMPATIBLE.ADVM` must be set to 11.2 or higher for the disk group. See ["Disk Group Compatibility Attributes"](#) on page 4-33.

The ALTER DISKGROUP VOLUME SQL statements enable you to manage Oracle ADVM volumes, including the functionality to add, resize, disable, enable, and drop volumes. For example:

```
SQL> ALTER DISKGROUP data ADD VOLUME volume1 SIZE 10G;
Diskgroup altered.
```

```
SQL> ALTER DISKGROUP data RESIZE VOLUME volume1 SIZE 15G;
Diskgroup altered.
```

```
SQL> ALTER DISKGROUP data DISABLE VOLUME volume1;
Diskgroup altered.
```

```
SQL> ALTER DISKGROUP data ENABLE VOLUME volume1;
Diskgroup altered.
```

```
SQL> ALTER DISKGROUP ALL DISABLE VOLUME ALL;
Diskgroup altered.
```

```
SQL> ALTER DISKGROUP data DROP VOLUME volume1;
Diskgroup altered.
```

If the volume is hosting an Oracle ACFS file system, then you cannot resize that volume with the SQL ALTER DISKGROUP statement. Instead you must use the `acfsutil size` command. For information, see ["acfsutil size"](#) on page 13-29.

For information about Oracle ADVM, see ["Overview of Oracle ASM Dynamic Volume Manager"](#) on page 5-10. For information about managing Oracle ADVM volumes with ASMCMD, see ["ASMCMD Volume Management Commands"](#) on page 12-58. For information about managing Oracle ADVM volumes with ASMCA, see ["Managing Oracle ADVM Volumes With Oracle ASM Configuration Assistant"](#) on page 11-9. For information about managing Oracle ADVM volumes with Oracle Enterprise Manager, see ["Managing Oracle ACFS with Oracle Enterprise Manager"](#) on page 10-1.

See Also: The ALTER DISKGROUP SQL statement in the *Oracle Database SQL Language Reference*

Dropping Disks from Disk Groups

To drop disks from a disk group, use the DROP DISK clause of the ALTER DISKGROUP statement. You can also drop all of the disks in specified failure groups using the DROP DISKS IN FAILGROUP clause.

When a disk is dropped, the disk group is rebalanced by moving all of the file extents from the dropped disk to other disks in the disk group. A drop disk operation might fail if not enough space is available on the other disks. The best approach is to perform both the add and drop operation with the same ALTER DISKGROUP statement. This has the benefit of rebalancing data extents only once and ensuring that there is enough space for the rebalance operation to succeed.

Caution: The ALTER DISKGROUP . . . DROP DISK statement returns before the drop and rebalance operations are complete. Do not reuse, remove, or disconnect the dropped disk until the HEADER_STATUS column for this disk in the V\$ASM_DISK view changes to FORMER. You can query the V\$ASM_OPERATION view to determine the amount of time remaining for the drop/rebalance operation to complete. For more information, refer to the *Oracle Database SQL Language Reference* and the *Oracle Database Reference*.

If you specify the `FORCE` clause for the drop operation, the disk is dropped even if Oracle ASM cannot read or write to the disk. You cannot use the `FORCE` flag when dropping a disk from an external redundancy disk group.

Caution: A `DROP FORCE` operation leaves data at reduced redundancy until the subsequent rebalance operation completes. This increases your exposure to data loss if there is a subsequent disk failure during rebalancing. `DROP FORCE` should be used only with great care.

You can also drop disks from a disk group with Oracle Enterprise Manager. See ["Dropping Disks from Disk Groups"](#) on page 9-11.

Example: Dropping Disks from Disk Groups

The statements in this example demonstrate how to drop disks from the disk group `data1` described in ["Example: Adding Disks to a Disk Group"](#) on page 4-12.

The following example drops `diska5` from disk group `data1`.

```
ALTER DISKGROUP data1 DROP DISK diska5;
```

The following example drops `diska5` from disk group `data1`, and also illustrates how multiple actions are possible with one `ALTER DISKGROUP` statement.

```
ALTER DISKGROUP data1 DROP DISK diska5
      ADD FAILGROUP failgrp1 DISK '/devices/diska9' NAME diska9;
```

Intelligent Data Placement

Intelligent Data Placement enables you to specify disk regions on Oracle ASM disks for best performance. Using the disk region settings, you can ensure that frequently accessed data is placed on the outermost (hot) tracks which have greater speed and higher bandwidth. In addition, files with similar access patterns are located physically close together, reducing latency. Intelligent Data Placement also allows the placement of primary and mirror extents into different hot or cold regions.

Intelligent Data 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.

Intelligent Data Placement works best for the following:

- Databases with data files that are accessed at different rates. A database that accesses all data files in the same way is unlikely to benefit from Intelligent Data Placement.
- Disk groups that are more than 25% full. If the disk group is only 25% full, the management overhead is unlikely to be worth any benefit.
- Disks that have better performance at the beginning of the media relative to the end. Because Intelligent Data Placement leverages the geometry of the disk, it is well suited to JBOD (just a bunch of disks). In contrast, a storage array with LUNs composed of concatenated volumes masks the geometry from Oracle ASM.

The `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` disk group attributes must be set to 11.2 or higher to use Intelligent Data Placement.

Intelligent Data Placement can be managed with the `ALTER DISKGROUP ADD` or `MODIFY TEMPLATE SQL` statements and the `ALTER DISKGROUP MODIFY FILE SQL` statement.

- The `ALTER DISKGROUP TEMPLATE SQL` statement includes a disk region clause for setting *hot/mirrorhot* or *cold/mirrorcold* regions in a template:

```
ALTER DISKGROUP data ADD TEMPLATE datafile_hot
  ATTRIBUTES (
    HOT
    MIRRORHOT);
```

- The `ALTER DISKGROUP ... MODIFY FILE SQL` statement that sets disk region attributes for *hot/mirrorhot* or *cold/mirrorcold* regions:

```
ALTER DISKGROUP data MODIFY FILE '+data/orcl/datafile/users.259.679156903'
  ATTRIBUTES (
    HOT
    MIRRORHOT);
```

When you modify the disk region settings for a file, this action applies to new extensions of the file, but existing file contents are not affected until a rebalance operation. To apply the new Intelligent Data Placement policy for existing file contents, you can manually initiate a rebalance. A rebalance operation uses the last specified policy for the file extents. For information on the rebalance operation, see ["Manually Rebalancing Disk Groups"](#) on page 4-17.

Oracle ASM Configuration Assistant (ASMCA) supports Intelligent Data Placement with template creation during disk group alterations. See ["Managing Disk Groups With Oracle ASM Configuration Assistant"](#) on page 11-6.

Oracle Enterprise Manager supports Intelligent Data Placement from the **Templates** page launched from the disk group page. See ["Managing Disk Group Templates with Oracle Enterprise Manager"](#) on page 9-20.

For information about Intelligent Data Placement details in views, see ["Viewing Disk Region Information"](#) on page 6-5.

Resizing Disks in Disk Groups

The `RESIZE` clause of `ALTER DISKGROUP` enables you to perform the following operations:

- Resize all disks in the disk group
- Resize specific disks
- Resize all of the disks in a specified failure group

If you do not specify a new size in the `SIZE` clause then Oracle ASM uses the size of the disk as returned by the operating system. The new size is written to the Oracle ASM disk header and if the size of the disk is increasing, then the new space is immediately available for allocation. If the size is decreasing, rebalancing must relocate file extents beyond the new size limit to available space below the limit. If the rebalance operation can successfully relocate all extents, then the new size is made permanent, otherwise the rebalance fails.

Example: Resizing Disks in Disk Groups

The following example resizes all of the disks in failure group `failgrp1` of disk group `data1`. If the new size is greater than disk capacity, the statement fails.

```
ALTER DISKGROUP data1
  RESIZE DISKS IN FAILGROUP failgrp1 SIZE 100G;
```

Undropping Disks in Disk Groups

The `UNDROP DISKS` clause of the `ALTER DISKGROUP` statement enables you to cancel all pending drops of disks within disk groups. If a drop disk operation has completed, then this statement cannot be used to restore it. This statement cannot be used to restore disks that are being dropped as the result of a `DROP DISKGROUP` statement, or for disks that are being dropped using the `FORCE` clause.

Example: Undropping Disks in Disk Groups

The following example cancels the dropping of disks from disk group `data1`:

```
ALTER DISKGROUP data1 UNDROP DISKS;
```

Manually Rebalancing Disk Groups

You can manually rebalance the files in a disk group using the `REBALANCE` clause of the `ALTER DISKGROUP` statement. This would normally not be required, because Oracle ASM automatically rebalances disk groups when their configuration changes. You might want to do a manual rebalance operation if you want to control the speed of what would otherwise be an automatic rebalance operation.

The `POWER` clause of the `ALTER DISKGROUP . . . REBALANCE` statement specifies the degree of parallelism, and thus the speed of the rebalance operation. It can be set to a value from 0 to 11. A value of 0 halts a rebalancing operation until the statement is either implicitly or explicitly re-run. The default rebalance power is set by the `ASM_POWER_LIMIT` initialization parameter. See "[Tuning Rebalance Operations](#)" on page 4-18 for more information.

The power level of an ongoing rebalance operation can be changed by entering the rebalance statement with a new level.

The `ALTER DISKGROUP . . . REBALANCE` command by default returns immediately so that you can issue other commands while the rebalance operation takes place asynchronously in the background. You can query the `V$ASM_OPERATION` view for the status of the rebalance operation.

If you want the `ALTER DISKGROUP . . . REBALANCE` command to wait until the rebalance operation is complete before returning, you can add the `WAIT` keyword to the `REBALANCE` clause. This is especially useful in scripts. The command also accepts a `NOWAIT` keyword, which invokes the default behavior of conducting the rebalance operation asynchronously. You can interrupt a rebalance running in wait mode by typing `CTRL-C` on most platforms. This causes the command to return immediately with the message `ORA-01013: user requested cancel of current operation`, and then to continue the rebalance operation asynchronously.

Additional rules for the rebalance operation include the following:

- An ongoing rebalance command is restarted if the storage configuration changes either when you alter the configuration, or if the configuration changes due to a failure or an outage. Furthermore, if the new rebalance fails because of a user error, then a manual rebalance may be required.
- The `ALTER DISKGROUP . . . REBALANCE` statement runs on a single node even if you are using Oracle Real Application Clusters (Oracle RAC).

- Oracle ASM can perform one disk group rebalance at a time on a given instance. Therefore, if you have initiated multiple rebalances on different disk groups, then Oracle processes this operation serially. However, you can initiate rebalances on different disk groups on different nodes in parallel.
- Rebalancing continues across a failure of the Oracle ASM instance performing the rebalance.
- The REBALANCE clause (with its associated POWER and WAIT/NOWAIT keywords) can also be used in ALTER DISKGROUP commands that add, drop, or resize disks.

Note: Oracle restarts the processing of an ongoing rebalance operation if the storage configuration changes. If the next rebalance operation fails because of a user error, then you may need to perform a manual rebalance.

Example: Manually Rebalancing a Disk Group

The following example manually rebalances the disk group `data2`. The command does not return until the rebalance operation is complete.

```
ALTER DISKGROUP data2 REBALANCE POWER 5 WAIT;
```

For more information about rebalancing operations, refer to "[Tuning Rebalance Operations](#)" on page 4-18.

Tuning Rebalance Operations

If the POWER clause is not specified in an ALTER DISKGROUP statement, or when rebalance is implicitly run by adding or dropping a disk, then the rebalance power defaults to the value of the ASM_POWER_LIMIT initialization parameter. You can adjust the value of this parameter dynamically.

The higher the power limit, the more quickly a rebalance operation can complete. Rebalancing takes longer with lower power values, but consumes fewer processing and I/O resources which are shared by other applications, such as the database.

The default value of 1 minimizes disruption to other applications. The appropriate value is dependent on your hardware configuration, performance requirements, and availability requirements

If a rebalance is in progress because a disk is manually or automatically dropped, then increasing the power of the rebalance shortens the time frame during which redundant copies of that data on the dropped disk are reconstructed on other disks.

The V\$ASM_OPERATION view provides information for adjusting ASM_POWER_LIMIT and the resulting power of rebalance operations. The V\$ASM_OPERATION view also gives an estimate in the EST_MINUTES column of the amount of time remaining for the rebalance operation to complete. You can see the effect of changing the rebalance power by observing the change in the time estimate.

See Also: "[Manually Rebalancing Disk Groups](#)" on page 4-17 for more information.

Oracle ASM Disk Discovery

Disk discovery is the mechanism used to find the operating system names for disks Oracle ASM can access. It is used to find all the disks that comprise a disk group to be

mounted, the disks an administrator wants to add to a disk group, or the disks the administrator might consider adding to a disk group. This section contains the following topics:

- [How A Disk is Discovered](#)
- [Disk Discovery Rules](#)
- [Improving Disk Discovery Time](#)

For additional information about disk discovery and the `ASM_DISKSTRING` initialization parameter, refer to "[ASM_DISKSTRING](#)" on page 3-7.

How A Disk is Discovered

When an Oracle ASM instance is initialized, Oracle ASM discovers and examines the contents of all of the disks that are in the paths that you designated with values in the `ASM_DISKSTRING` initialization parameter.

Disk discovery also occurs when you:

- Run the following SQL statements
 - Mount a disk group with `ALTER DISKGROUP ... MOUNT`
 - Online a disk with `ALTER DISKGROUP ... ONLINE DISK`
 - Add a disk to a disk group with `CREATE` or `ALTER DISKGROUP...ADD DISK`
 - Resize a disk in a disk group with `ALTER DISKGROUP...RESIZE DISK`
 - Query with `SELECT ... FROM V$ASM_DISKGROUP` or `V$ASM_DISK` views
- Run Oracle Enterprise Manager or Oracle ASM Configuration Assistant (ASMCA) operations that invoke the SQL statements previously listed
- Run `ASMCMD` commands that perform the same operations as the SQL statements previously listed

After Oracle ASM successfully discovers a disk, the disk appears in the `V$ASM_DISK` view. Disks that belong to a disk group, that is, disks that have a disk group name in the disk header, show a header status of `MEMBER`. Disks that were discovered, but that have not yet been assigned to a disk group, have a status of either `CANDIDATE` or `PROVISIONED`. Disks that previously belonged to a disk group and were dropped cleanly from the disk group have a status of `FORMER`.

The `PROVISIONED` status implies that an additional platform-specific action has been taken by an administrator to make the disk available for Oracle ASM. For example, on Windows computers, the administrator might have used `asmtool` or `asmtoolg` to stamp the disk with a header. On Linux computers, the administrator might have used `ASMLIB` to prepare the disk for Oracle ASM.

[Example 4-6](#) shows a SQL query on `V$ASM_DISK` that displays the header status of a group of disks.

Example 4-6 Querying V\$ASM_DISK for Header Status

```
SQL> SELECT name, header_status, path FROM V$ASM_DISK
       WHERE path LIKE '/devices/disk0%'
```

NAME	HEADER_STATUS	PATH

	FORMER	/devices/disk02
	FORMER	/devices/disk01

```

          CANDIDATE      /devices/disk07
DISK06   MEMBER        /devices/disk06
DISK05   MEMBER        /devices/disk05
DISK04   MEMBER        /devices/disk04
DISK03   MEMBER        /devices/disk03
7 rows selected.

```

See Also: *Oracle Database Reference* for information about the header status of an Oracle ASM disk that is displayed in the `V$ASM_DISK` view

Disk Discovery Rules

The rules for discovering Oracle ASM disks are as follows:

- Oracle ASM can discover up to 10,000 disks. That is, if more than 10,000 disks match the `ASM_DISKSTRING` initialization parameter, then Oracle ASM discovers only the first 10,000.
- Oracle ASM only discovers disk partitions. Oracle ASM does not discover partitions that include the partition table.
- From the perspective of the installation, candidate disks are those that have the `CANDIDATE`, `PROVISIONED`, or `FORMER` header status. These disks with a `CANDIDATE`, `PROVISIONED`, or `FORMER` status can be added to Oracle ASM disk groups without using the `FORCE` flag.
- When adding a disk, the `FORCE` option must be used if Oracle ASM recognizes that the disk was managed by Oracle. Such a disk appears in the `V$ASM_DISK` view with a status of `FOREIGN`. In this case, you can only add the disk to a disk group by using the `FORCE` keyword.
- `MEMBER` disks can usually be added to a disk group by specifying the `FORCE` flag, provided that the disks are not part of a currently mounted disk group.

In addition, Oracle ASM identifies the following configuration errors during discovery:

- Multiple paths to the same disk

In this case, if the disk is part of a disk group, then disk group mount fails. If the disk is being added to a disk group with the `ADD DISK` or `CREATE DISKGROUP` command, then the command fails. To correct the error, adjust the `ASM_DISKSTRING` value so that Oracle ASM does not discover multiple paths to the same disk. Or if you are using multipathing software, then ensure that you include only the pseudo-device name in the `ASM_DISKSTRING` value. See "[Oracle ASM and Multipathing](#)" on page 2-3.
- Multiple Oracle ASM disks with the same disk header

This can be caused by having copied one disk onto another. In this case, the disk group mount operation fails.

Improving Disk Discovery Time

The value for the `ASM_DISKSTRING` initialization parameter is an operating system-dependent value that Oracle ASM uses to limit the set of paths that the discovery process uses to search for disks. When a new disk is added to a disk group, each Oracle ASM instance that has the disk group mounted must be able to discover the new disk using its `ASM_DISKSTRING`.

In many cases, the default value (NULL) is sufficient. Using a more restrictive value might reduce the time required for Oracle ASM to perform discovery, and thus improve disk group mount time or the time for adding a disk to a disk group. Oracle might need to dynamically change the `ASM_DISKSTRING` before adding a disk so that the new disk is discovered through this parameter.

The default value of `ASM_DISKSTRING` might not find all disks in all situations. If your site is using a third-party vendor `ASMLIB`, then the vendor might have discovery string conventions that you must use for `ASM_DISKSTRING`. In addition, if your installation uses multipathing software, then the software might place pseudo-devices in a path that is different from the operating system default. See "[Oracle ASM and Multipathing](#)" on page 2-3 and consult the multipathing vendor documentation for details.

Managing Capacity in Disk Groups

When Oracle ASM provides redundancy, such as when you create a disk group with `NORMAL` or `HIGH` redundancy, you must have sufficient capacity in each disk group to manage a re-creation of data that is lost after a failure of one or two failure groups. After one or more disks fail, the process of restoring redundancy for all data requires space from the surviving disks in the disk group. If not enough space remains, then some files might end up with reduced redundancy.

Reduced redundancy means that one or more extents in the file are not mirrored at the expected level. For example, a reduced redundancy file in a high redundancy disk group has at least one file extent with two or fewer total copies of the extent instead of three. In the case of unprotected files, data extents could be missing altogether. Other causes of reduced redundancy files are disks running out of space or an insufficient number of failure groups. The `REDUNDANCY_LOWERED` column in the `V$ASM_FILE` view provides information about files with reduced redundancy.

The following guidelines help ensure that you have sufficient space to restore full redundancy for all disk group data after the failure of one or more disks.

- Normal redundancy disk group - It is best to have enough free space in your disk group to tolerate the loss of all disks in one failure group. The amount of free space should be equivalent to the size of the largest failure group.
- High redundancy disk group - It is best to have enough free space to cope with the loss of all disks in two failure groups. The amount of free space should be equivalent to the sum of the sizes of the two largest failure groups.

Note: When you lose multiple disks from multiple failure groups, then you could lose both the primary and the redundant copies of your data. In addition, if you do not have enough capacity to restore redundancy, then Oracle ASM can continue to operate. However, if another disk fails, then the system may not be able to tolerate additional failures.

The `V$ASM_DISKGROUP` view contains the following columns that contain information to help you manage capacity:

- `REQUIRED_MIRROR_FREE_MB` indicates the amount of space that must be available in a disk group to restore full redundancy after the worst failure that can be tolerated by the disk group without adding additional storage. This means that there are sufficient failure groups to restore redundancy. Also, this worst failure

refers to a permanent failure where the disks must be dropped, not the case where the disks go offline and then back online.

The amount of space displayed in this column takes the effects of mirroring into account. The value is computed as follows:

- Normal redundancy disk group with more than two failure groups
The value is the total raw space for all of the disks in the largest failure group. The largest failure group is the one with the largest total raw capacity. For example, if each disk is in its own failure group, then the value would be the size of the largest capacity disk.
- High redundancy disk group with more than three failure groups
The value is the total raw space for all of the disks in the two largest failure groups.
- `USABLE_FILE_MB` indicates the amount of free space, adjusted for mirroring, that is available for new files to restore redundancy after a disk failure. `USABLE_FILE_MB` is computed by subtracting `REQUIRED_MIRROR_FREE_MB` from the total free space in the disk group and then adjusting the value for mirroring. For example, in a normal redundancy disk group where by default the mirrored files use disk space equal to twice their size, if 4 GB of actual usable file space remains, then `USABLE_FILE_MB` equals roughly 2 GB. You can then add a file that is up to 2 GB in size.
- `TOTAL_MB` is the total usable capacity of a disk group in megabytes. The calculations for data in this column take the disk header overhead into consideration. The disk header overhead depends on the number of Oracle ASM disks and Oracle ASM files. This value is typically about 1% of the total raw storage capacity. For example, if the total LUN capacity provisioned for Oracle ASM is 100 GB, then the value in the `TOTAL_MB` column would be about 99 GB.
- `FREE_MB` is the unused capacity of the disk group in megabytes, without considering any data imbalance. Therefore, there may be situations where the value in the `FREE_MB` column shows unused capacity but because one of the Oracle ASM disks is full, database writes fail because of the imbalance in the disk group. Make sure that you initiate a manual rebalance to force even data distribution which results in an accurate presentation of the values in the `FREE_MB` column.

With fine grain striping using 128 KB, the storage is preallocated to be eight times the AU size. Therefore, the data file size may appear slightly larger on Oracle ASM than on a local file system because of the preallocation.

When you use Oracle ASM normal or high redundancy, the disk space utilization becomes more complex to measure because it depends on several variables.

Note: The values in the `TOTAL_MB` and `FREE_MB` columns best describe space usage when you do not configure Oracle ASM mirroring, that is, when you use external redundancy.

The results from the following query show capacity metrics for a normal redundancy disk group that consists of six 1 GB (1024 MB) disks, each in its own failure group:

```
SQL> SELECT name, type, total_mb, free_mb, required_mirror_free_mb,
         usable_file_mb FROM V$ASM_DISKGROUP;
```

NAME	TYPE	TOTAL_MB	FREE_MB	REQUIRED_MIRROR_FREE_MB	USABLE_FILE_MB
------	------	----------	---------	-------------------------	----------------

```

-----
DATA          NORMAL          6144          3768          1024          1372

```

The `REQUIRED_MIRROR_FREE_MB` column shows that 1 GB of extra capacity must be available to restore full redundancy after one or more disks fail. The first three numeric columns in the query results are raw numbers. That is, they do not take redundancy into account. Only the last column is adjusted for normal redundancy. In the query output example for the `data` disk group, the calculation is as follows:

$$\text{FREE_MB} - \text{REQUIRED_MIRROR_FREE_MB} = 2 * \text{USABLE_FILE_MB}$$

$$3768 - 1024 = 2 * 1372 = 2744$$

Negative Values of `USABLE_FILE_MB`

Due to the relationship between `FREE_MB`, `REQUIRED_MIRROR_FREE_MB`, and `USABLE_FILE_MB`, `USABLE_FILE_MB` can become negative. Although this is not necessarily a critical situation, it does mean that:

- Depending on the value of `FREE_MB`, you may not be able to create new files.
- The next failure might result in files with reduced redundancy.

If `USABLE_FILE_MB` becomes negative, it is strongly recommended that you add more space to the disk group as soon as possible.

Oracle ASM Mirroring and Disk Group Redundancy

This section contains the following topics:

- [Oracle ASM Mirroring and Failure Groups](#)
- [Oracle ASM Recovery from Read and Write I/O Errors](#)
- [Oracle ASM Fast Mirror Resync](#)
- [Preferred Read Failure Groups](#)

Oracle ASM Mirroring and Failure Groups

If you specify mirroring for a file, then Oracle ASM automatically stores redundant copies of the file extents in separate failure groups. Failure groups apply only to normal and high redundancy disk groups. You can define the failure groups for each disk group when you create or alter the disk group.

There are three types of disk groups based on the Oracle ASM redundancy level. [Table 4-1](#) lists the types with their supported and default mirroring levels. The default mirroring levels indicate the mirroring level with which each file is created unless a different mirroring level is designated.

Table 4-1 Mirroring Options for Oracle ASM Disk Group Types

Disk Group Type	Supported Mirroring Levels	Default Mirroring Level
External redundancy	Unprotected (none)	Unprotected
Normal redundancy	Two-way, three-way, unprotected (none)	Two-way
High redundancy	Three-way	Three-way

The redundancy level controls how many disk failures are tolerated without dismounting the disk group or losing data. Each file is allocated based on its own redundancy, but the default comes from the disk group.

The redundancy levels are:

- External redundancy
Oracle ASM does not provide mirroring redundancy and relies on the storage system to provide RAID functionality. Any write error cause a forced dismount of the disk group. All disks must be located to successfully mount the disk group.
- Normal redundancy
Oracle ASM provides two-way mirroring by default, which means that all files are mirrored so that there are two copies of every **extent**. A loss of one Oracle ASM disk is tolerated. You can optionally choose three-way or unprotected mirroring.
- High redundancy
Oracle ASM provides triple mirroring by default. A loss of two Oracle ASM disks in different failure groups is tolerated.

If there are not enough online failure groups to satisfy the file mirroring (redundancy attribute value) specified in the disk group file type template, Oracle ASM allocates as many mirrors copies as possible and subsequently allocates the remaining mirrors when sufficient online failure groups are available. For information about specifying Oracle ASM disk group templates, see "[Managing Disk Group Templates](#)" on page 7-15.

Failure groups enable the mirroring of metadata and user data. System reliability can diminish if your environment has an insufficient number of failure groups.

This section contains these topics:

- [Oracle ASM Failure Groups](#)
- [How Oracle ASM Manages Disk Failures](#)
- [Guidelines for Using Failure Groups](#)
- [Failure Group Frequently Asked Questions](#)

Oracle ASM Failure Groups

Failure groups are used to store mirror copies of data. When Oracle ASM allocates an extent for a normal redundancy file, Oracle ASM allocates a primary copy and a secondary copy. Oracle ASM chooses the disk on which to store the secondary copy so that it is in a different failure group than the primary copy. Each copy is on a disk in a different failure group so that the simultaneous failure of all disks in a failure group does not result in data loss.

A failure group is a subset of the disks in a disk group, which could fail at the same time because they share hardware. The failure of common hardware must be tolerated. Four drives that are in a single removable tray of a large JBOD (Just a Bunch of Disks) array should be in the same failure group because the tray could be removed making all four drives fail at the same time. Drives in the same cabinet could be in multiple failure groups if the cabinet has redundant power and cooling so that it is not necessary to protect against failure of the entire cabinet. However, Oracle ASM mirroring is not intended to protect against a fire in the computer room that destroys the entire cabinet.

There are always failure groups even if they are not explicitly created. If you do not specify a failure group for a disk, then Oracle automatically creates a new failure

group containing just that disk, except for disk groups containing disks on Oracle Exadata cells.

A normal redundancy disk group must contain at least two failure groups. A high redundancy disk group must contain at least three failure groups. However, Oracle recommends using several failure groups. A small number of failure groups, or failure groups of uneven capacity, can create allocation problems that prevent full use of all of the available storage.

Failure groups can be specified as regular or quorum failure groups. For information about quorum failure groups, see "[Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#)" on page 4-10.

See Also: *Oracle Exadata documentation* for information about Oracle Exadata failure groups

How Oracle ASM Manages Disk Failures

Depending on the redundancy level of a disk group and how you define failure groups, the failure of one or more disks could result in either of the following:

- The disks are first taken offline and then automatically dropped. In this case, the disk group remains mounted and serviceable. In addition, because of mirroring, all of the disk group data remains accessible. After the disk drop operation, Oracle ASM performs a rebalance to restore full redundancy for the data on the failed disks.
- The entire disk group is automatically dismounted, which means loss of data accessibility.

Guidelines for Using Failure Groups

The following are guidelines for using failure groups:

- Each disk in a disk group can belong to only one failure group.
- Failure groups should all be of the same size. Failure groups of different sizes may lead to reduced availability.
- Oracle ASM requires at least two failure groups to create a normal redundancy disk group and at least three failure groups to create a high redundancy disk group.

Failure Group Frequently Asked Questions

This section discusses frequently asked questions about failure group under the following topics:

- [How Many Failure Groups Should I Create?](#)
- [How are Multiple Failure Groups Recovered after Simultaneous Failures?](#)
- [When Should External, Normal, or High Redundancy Be Used?](#)

How Many Failure Groups Should I Create?

Choosing the number of failure groups to create depends on the types of failures that need to be tolerated without data loss. For small numbers of disks, such as fewer than 20, it is usually best to use the default failure group creation that puts every disk in its own failure group.

Using the default failure group creation for small numbers of disks is also applicable for large numbers of disks where your main concern is disk failure. For example, a

disk group might be configured from several small modular disk arrays. If the system must continue operating when an entire modular array fails, then a failure group should consist of all of the disks in one module. If one module fails, then all of the data on that module is relocated to other modules to restore redundancy. Disks should be placed in the same failure group if they depend on a common piece of hardware whose failure must be tolerated with no loss of availability.

How are Multiple Failure Groups Recovered after Simultaneous Failures?

A simultaneous failure can occur if there is a failure of a piece of hardware used by multiple failure groups. This type of failure usually forces a dismount of the disk group if all disks are unavailable.

When Should External, Normal, or High Redundancy Be Used?

Oracle ASM mirroring runs on the database server and Oracle recommends to off load this processing to the storage hardware RAID controller by using external redundancy. You can use normal redundancy in the following scenarios:

- Storage system does not have RAID controller
- Mirroring across storage arrays
- Extended cluster configurations

In general, Oracle ASM mirroring is the Oracle alternative to third party logical volume managers. Oracle ASM mirroring eliminates the need to deploy additional layers of software complexity in your Oracle database environment.

Oracle ASM Recovery from Read and Write I/O Errors

Read errors can be the result of a loss of access to the entire disk or media corruptions on an otherwise a healthy disk. Oracle ASM tries to recover from read errors on corrupted sectors on a disk. When a read error by the database or Oracle ASM triggers the Oracle ASM instance to attempt bad block remapping, Oracle ASM reads a good copy of the extent and copies it to the disk that had the read error.

- If the write to the same location succeeds, then the underlying allocation unit (sector) is deemed healthy. This might be because the underlying disk did its own bad block reallocation.
- If the write fails, Oracle ASM attempts to write the extent to a new allocation unit on the same disk. If this write succeeds, the original allocation unit is marked as unusable. If the write fails, the disk is taken offline.

One unique benefit on Oracle ASM-based mirroring is that the database instance is aware of the mirroring. For many types of logical corruptions such as a bad checksum or incorrect System Change Number (SCN), the database instance proceeds through the mirror side looking for valid content and proceeds without errors. If the process in the database that encountered the read is in a position to obtain the appropriate locks to ensure data consistency, it writes the correct data to all mirror sides.

When encountering a write error, a database instance sends the Oracle ASM instance a *disk offline* message.

- If database can successfully complete a write to at least one extent copy and receive acknowledgment of the offline disk from Oracle ASM, the write is considered successful.
- If the write to all mirror side fails, database takes the appropriate actions in response to a write error such as taking the tablespace offline.

When the Oracle ASM instance receives a write error message from a database instance or when an Oracle ASM instance encounters a write error itself, Oracle ASM instance attempts to take the disk offline. Oracle ASM consults the Partner Status Table (PST) to see whether any of the disk's partners are offline. If too many partners are already offline, Oracle ASM forces the dismounting of the disk group. Otherwise, Oracle ASM takes the disk offline.

The ASMCMD `remap` command was introduced to address situations where a range of bad sectors exists on a disk and must be corrected before Oracle ASM or database I/O. For information about the `remap` command, see ["remap"](#) on page 12-45.

Oracle ASM Fast Mirror Resync

Restoring the redundancy of an Oracle ASM disk group after a transient disk path failure can be time consuming. This is especially true if the recovery process requires rebuilding an entire Oracle ASM failure group. Oracle ASM fast mirror resync significantly reduces the time to resynchronize a failed disk in such situations. When you replace the failed disk, Oracle ASM can quickly resynchronize the Oracle ASM disk extents.

Note: To use this feature, the disk group compatibility attributes must be set to 11.1 or higher. For more information, refer to ["Disk Group Compatibility"](#) on page 4-31.

Any problems that make a failure group temporarily unavailable are considered transient failures that can be recovered by the Oracle ASM fast mirror resync feature. For example, transient failures can be caused by disk path malfunctions, such as cable failures, host bus adapter failures, controller failures, or disk power supply interruptions.

Oracle ASM fast resync keeps track of pending changes to extents on an OFFLINE disk during an outage. The extents are resynced when the disk is brought back online.

By default, Oracle ASM drops a disk in 3.6 hours after it is taken offline. You can set the `DISK_REPAIR_TIME` disk group attribute to delay the drop operation by specifying a time interval to repair the disk and bring it back online. The time can be specified in units of minutes (m or M) or hours (h or H). If you omit the unit, then the default unit is hours. The `DISK_REPAIR_TIME` disk group attribute can only be set with the `ALTER DISKGROUP` SQL statement.

If the attribute is not set explicitly, then the default value (3.6h) applies to disks that have been set to OFFLINE mode without an explicit `DROP AFTER` clause. Disks taken offline due to I/O errors do not have a `DROP AFTER` clause.

The default `DISK_REPAIR_TIME` attribute value is an estimate that should be adequate for most environments. However, make sure that the attribute value is set to the amount of time that you think is necessary in your environment to fix any transient disk error and that you are willing to tolerate reduced data redundancy.

The elapsed time (since the disk was set to OFFLINE mode) is incremented only when the disk group containing the offline disks is mounted. The `REPAIR_TIME` column of `V$ASM_DISK` shows the amount of time left (in seconds) before an offline disk is dropped. After the specified time has elapsed, Oracle ASM drops the disk. You can override this attribute with an `ALTER DISKGROUP DISK OFFLINE` statement and the `DROP AFTER` clause.

Note: If a disk is offlined by Oracle ASM because of an I/O (write) error or is explicitly offlined using the ALTER DISKGROUP... OFFLINE statement without the DROP AFTER clause, then the value that is specified for the DISK_REPAIR_TIME attribute for the disk group is used. If this attribute value is changed with the ALTER DISKGROUP... SET ATTRIBUTE 'disk_repair_time' statement before this offlined disk is dropped, then the new (current) value of the attribute is used by the Oracle ASM disk offline functionality but the elapsed time is not reset. You can confirm this behavior by viewing the Oracle ASM alert log.

If an offline disk is taken offline for a second time, then the elapsed time is reset and restarted. If another time is specified with the DROP AFTER clause for this disk, the first value is overridden and the new value applies. A disk that is in OFFLINE mode cannot be dropped with an ALTER DISKGROUP DROP DISK statement; an error is returned if attempted. If for some reason the disk must be dropped (such as the disk cannot be repaired) before the repair time has expired, a disk can be dropped immediately by issuing a second OFFLINE statement with a DROP AFTER clause specifying 0h or 0m.

You can use ALTER DISKGROUP to set the DISK_REPAIR_TIME attribute to a specified hour or minute value, such as 4.5 hours or 270 minutes. For example:

```
ALTER DISKGROUP data SET ATTRIBUTE 'disk_repair_time' = '4.5h'  
ALTER DISKGROUP data SET ATTRIBUTE 'disk_repair_time' = '270m'
```

After you repair the disk, run the SQL statement ALTER DISKGROUP DISK ONLINE. This statement brings a repaired disk group back online to enable writes so that no new writes are missed. This statement also starts a procedure to copy of all of the extents that are marked as stale on their redundant copies.

If a disk goes offline when the Oracle ASM instance is in rolling upgrade mode, the disk remains offline until the rolling upgrade has ended and the timer for dropping the disk is stopped until the Oracle ASM cluster is out of rolling upgrade mode. See ["Using Oracle ASM Rolling Upgrade"](#) on page 3-21. Examples of taking disks offline and bringing them online follow.

The following example takes disk DATA_001 offline and drops it after five minutes.

```
ALTER DISKGROUP data OFFLINE DISK DATA_001 DROP AFTER 5m;
```

The next example takes the disk DATA_001 offline and drops it after the time period designated by DISK_REPAIR_TIME elapses:

```
ALTER DISKGROUP data OFFLINE DISK DATA_001;
```

This example takes all of the disk in failure group FG2 offline and drops them after the time period designated by DISK_REPAIR_TIME elapses. IF you used a DROP AFTER clause, then the disks would be dropped after the specified time:

```
ALTER DISKGROUP data OFFLINE DISKS IN FAILGROUP FG2;
```

The next example brings all of the disks in failure group FG2 online:

```
ALTER DISKGROUP data ONLINE DISKS IN FAILGROUP FG2;
```

This example brings only disk DATA_001 online:

```
ALTER DISKGROUP data ONLINE DISK DATA_001;
```


This example brings all of the disks in disk group DATA online:

```
ALTER DISKGROUP data ONLINE ALL;
```

Querying the V\$ASM_OPERATION view while you are running any of these types of ALTER DISKGROUP ... ONLINE statements displays the name and state of the current operation that you are performing. For example, the query:

```
SELECT GROUP_NUMBER, OPERATION, STATE FROM V$ASM_OPERATION;
```

Displays output similar to the following:

```
GROUP_NUMBER OPERA STAT
-----
1 ONLIN RUN
```

An OFFLINE operation is not displayed in a V\$ASM_OPERATION view query.

See Also: *Oracle Database SQL Language Reference* for information about ALTER DISKGROUP and CREATE DISKGROUP

Preferred Read Failure Groups

When you configure Oracle ASM failure groups, it might be more efficient for a node to read from an extent that is closest to the node, even if that extent is a secondary extent. In other words, you can configure Oracle ASM to read from a secondary extent if that extent is closer to the node instead of Oracle ASM reading from the primary copy which might be farther from the node. Using preferred read failure groups is most useful in extended clusters.

To use this feature, Oracle recommends that you configure at least one mirrored extent copy from a disk that is local to a node in an extended cluster. However, a failure group that is preferred for one instance might be remote to another instance in the same Oracle RAC database. The parameter setting for preferred read failure groups is instance specific.

Both the Oracle ASM clients and Oracle ASM require Oracle Database 11g release 1 (11.1) or higher to use preferred read failure groups.

Note: By default, when you create a disk group, every disk in the disk group belongs to one failure group. Oracle does not recommend that you configure more than one preferred read failure group for each instance in a disk group. If you configure more than one preferred read failure group for each instance, then Oracle writes messages to an alert log.

See Also: *Oracle Real Application Clusters Administration and Deployment Guide* for information about configuring preferred read disks in extended clusters

Configuring and Administering Preferred Read Failure Groups

To configure this feature, set the ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter to specify a list of failure group names as preferred read disks. For more information about this initialization parameter, refer to "[ASM_PREFERRED_READ_FAILURE_GROUPS](#)" on page 3-8.

Set the parameter where *diskgroup* is the name of the disk group and *failuregroup* is the name of the failure group, separating these variables with a period. Oracle ASM ignores the name of a failure group that you use in this parameter setting if the failure group does not exist in the named disk group. You can append multiple values using commas as a separator as follows:

```
ASM_PREFERRED_READ_FAILURE_GROUPS = diskgroup.failuregroup,...
```

In an extended cluster, the failure groups that you specify with settings for the `ASM_PREFERRED_READ_FAILURE_GROUPS` parameter should only contain disks that are local to the instance. For normal redundancy disk groups, there should be only one failure group on each site of the extended cluster.

If there is more than one mirrored copy and you have set a value for the `ASM_PREFERRED_READ_FAILURE_GROUPS` parameter, then Oracle ASM first reads the copy that resides on a preferred read disk. If that read fails, then Oracle ASM attempts to read from the next mirrored copy that might not be on a preferred read disk.

Having more than one failure group on one site can cause the loss of access to the disk group by the other sites if the site containing more than one failure group fails. In addition, by having more than one failure group on a site, an extent might not be mirrored to another site. This can diminish the read performance of the failure group on the other site.

For example, for a normal redundancy disk group, if a site contains two failure groups of a disk group, then Oracle ASM might put both mirror copies of an extent on the same site. In this configuration, Oracle ASM cannot protect against data loss from a site failure.

You should configure at most two failure groups on a site for a high redundancy disk group. If there are three sites in an extended cluster, for the same reason previously mentioned, then you should only create one failure group.

For a two-site extended cluster, a normal redundancy disk group only has two failure groups. In this case, you can only specify one failure group as a preferred read failure group for each instance.

You can use views to identify preferred read failure groups, such as the `V$ASM_DISK` view that shows whether a disk is a preferred read disk by the value in the `PREFERRED_READ` column. You can also use `V$ASM_DISK` to verify whether local disks in an extended cluster are preferred read disks. Use the Oracle ASM disk I/O statistics to verify that read operations are using the preferred read disks that you configured.

If a disk group is not optimally configured for an extended cluster, then Oracle ASM records warning messages in the alert logs. To identify specific performance issues with Oracle ASM preferred read failure groups, use the `V$ASM_DISK_IOSTAT` view. This view displays disk I/O statistics for each Oracle ASM client. You can also query the `V$ASM_DISK_IOSTAT` view on a database instance. However, this query only shows the I/O statistics for the database instance. In general, optimal preferred read extended cluster configurations balance performance with disk group availability.

See Also: *Oracle Database Reference* for details about the `V$ASM*` dynamic performance views

Performance and Scalability Considerations for Disk Groups

This section discusses the following considerations for evaluating disk group performance:

- [Determining the Number of Disk Groups](#)
- [Performance Characteristics When Grouping Disks](#)
- [Oracle ASM Storage Limits](#)

Determining the Number of Disk Groups

Use the following criteria to determine the number of disk groups to create:

- Disks in a given disk group should have similar size and performance characteristics. If you have several different types of disks in terms of size and performance, then create several disk groups that contain similar characteristics.
- Create separate disk groups for your database files and fast recovery area for backup files. This configuration allows fast recovery in case of a disk group failure.

Performance Characteristics When Grouping Disks

Oracle ASM load balances the file activity by uniformly distributing file extents across all of the disks in a disk group. For this technique to be effective it is important that disks in a disk group be of similar performance characteristics. For example, the newest and fastest disks might reside in a disk group reserved for the database work area, and slower drives could reside in a disk group reserved for the fast recovery area.

There might be situations where it is acceptable to temporarily have disks of different sizes and performance co-existing in a disk group. This would be the case when migrating from an old set of disks to a new set of disks. The new disks would be added and the old disks dropped. As the old disks are dropped, their storage is migrated to the new disks while the disk group is online.

Oracle ASM Storage Limits

Oracle ASM has the following limits:

- 63 disk groups in a storage system
- 10,000 Oracle ASM disks in a storage system
- 2 terabytes (TB) maximum storage for each Oracle ASM disk not on Oracle Exadata Storage
- 20 petabytes (PB) maximum storage not on Oracle Exadata Storage
- 4 PB maximum storage for each Oracle ASM disk on Oracle Exadata Storage
- 40 exabyte (EB) maximum storage for Oracle Exadata Storage
- 1 million files for each disk group

File size limits are dependent on the setting of disk group compatibility attributes. Oracle Database supports data file sizes up to 128 TB when the `COMPATIBLE . RDBMS` disk group attribute is set to 11.1 or greater. Oracle ASM supports file sizes greater than 128 TB in any redundancy mode. This provides near unlimited capacity for future growth. For information about Oracle ASM disk group compatibility attributes, see ["Disk Group Compatibility"](#) on page 4-31. For information about Oracle ASM file size limits, see [Table 4-4](#) on page 4-38.

Disk Group Compatibility

This section describes disk group compatibility under the following topics:

- [Overview of Disk Group Compatibility](#)
- [Disk Group Compatibility Attributes](#)
- [Setting Disk Group Compatibility Attributes](#)
- [Considerations When Setting Disk Group Compatibility in Replicated Environments](#)

Overview of Disk Group Compatibility

Advancing the disk group compatibility settings enables you to use the new Oracle ASM features that are available in a later release. For example, a disk group with the disk group compatibility attributes set to 11.2 can take advantage of new Oracle 11g release 2 (11.2) features, such as Oracle ASM volumes in disk groups and Oracle ASM File Access Control. See [Table 4-3](#) on page 4-36 for the features enabled for combinations of compatibility attribute settings.

The disk group compatibility feature also enables environments to interoperate when they use disk groups from both Oracle Database 10g and Oracle Database 11g. For example, disk group compatibility settings that are set to Oracle Database 10g enable a 10g database client to access a disk group created with Oracle ASM 11g.

The disk group attributes that determine compatibility are `COMPATIBLE.ASM`, `COMPATIBLE.RDBMS`, and `COMPATIBLE.ADVM`. The `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` attribute settings determine the minimum Oracle Database software version numbers that a system can use for Oracle ASM and the database instance types respectively. For example, if the Oracle ASM compatibility setting is 11.2, and RDBMS compatibility is set to 11.1, then the Oracle ASM software version must be at least 11.2, and the Oracle Database client software version must be at least 11.1. The `COMPATIBLE.ADVM` attribute determines where the Oracle ASM Dynamic Volume Manager feature can be used to create Oracle ASM volumes in disk groups.

When you create a disk group, you can specify the disk group compatibility attribute settings in the `CREATE DISKGROUP` SQL statement. The `ALTER DISKGROUP` SQL statement can update the compatible attribute settings for existing disk groups. If not specified when using the `CREATE DISKGROUP` SQL statement, 10.1 is the default setting for both the `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` attributes for Oracle ASM in Oracle Database 11g. The `COMPATIBLE.ADVM` attribute is empty if it is not set. See [Table 4-2](#) on page 4-34 for examples of valid combinations of compatible attribute settings.

Notes:

- The disk group compatibility settings determine whether your environment can use the latest Oracle ASM features.
 - The disk group compatibility settings can only be advanced; you cannot revert to a lower compatibility setting. See ["Reverting Disk Group Compatibility"](#) on page 4-37.
 - The `COMPATIBLE.ASM` attribute must be advanced before advancing other disk group compatibility attributes and its value must be greater than or equal to the value of other disk group compatibility attributes.
-
-

Disk Group Compatibility Attributes

The disk group compatibility attributes specify the disk group compatibility settings for Oracle ASM and database instances. These attributes are described under the following topics:

- [COMPATIBLE.ASM](#)
- [COMPATIBLE.RDBMS](#)
- [COMPATIBLE.ADVM](#)

COMPATIBLE.ASM

The value for the disk group `COMPATIBLE.ASM` attribute determines the minimum software version for an Oracle ASM instance that can use the disk group. This setting also affects the format of the data structures for the Oracle ASM metadata on the disk. The format of other file contents is determined by Oracle ASM Dynamic Volume Manager (Oracle ADVM) and the database instance.

For Oracle ASM in Oracle Database 11g, 10.1 is the default setting for the `COMPATIBLE.ASM` attribute when using the SQL `CREATE DISKGROUP` statement, the ASMCMD `mkdgm` command, and Oracle Enterprise Manager **Create Disk Group** page. When creating a disk group with ASMCA, the default setting is 11.2.

COMPATIBLE.RDBMS

The value for the disk group `COMPATIBLE.RDBMS` attribute determines the minimum `COMPATIBLE` database initialization parameter setting for any database instance that is allowed to use the disk group. Before advancing the `COMPATIBLE.RDBMS` attribute, ensure that the values for the `COMPATIBLE` initialization parameter for all of the databases that access the disk group are set to at least the value of the new setting for `COMPATIBLE.RDBMS`.

For example, if the `COMPATIBLE` initialization parameters of the databases are set to either 11.1 or 11.2, then `COMPATIBLE.RDBMS` can be set to any value between 10.1 and 11.1 inclusively.

For Oracle ASM in Oracle Database 11g, 10.1 is the default setting for the `COMPATIBLE.RDBMS` attribute when using the SQL `CREATE DISKGROUP` statement, the ASMCMD `mkdgm` command, ASMCA **Create Disk Group** page, and Oracle Enterprise Manager **Create Disk Group** page.

Note: The database initialization parameter `COMPATIBLE` enables you to use a new release of Oracle Database, while at the same time guaranteeing backward compatibility with an earlier release. See *Oracle Database Reference* for more information about the `COMPATIBLE` initialization parameter.

COMPATIBLE.ADVM

The value for the disk group `COMPATIBLE.ADVM` attribute determines whether the disk group can contain Oracle ASM volumes. The value must be set to 11.2 or higher. Before setting this attribute, the `COMPATIBLE.ASM` value must be 11.2 or higher. Also, the Oracle ADVM volume drivers must be loaded.

By default, the value of the `COMPATIBLE.ADVM` attribute is empty until set.

For more information about Oracle ADVM, see "[Overview of Oracle ASM Dynamic Volume Manager](#)" on page 5-10.

Setting Disk Group Compatibility Attributes

This section discusses the settings of the disk group compatibility attributes and how to set the attribute values with the `CREATE DISKGROUP` or `ALTER DISKGROUP` SQL statement.

This section contains these topics:

- [Valid Combinations of Compatibility Attribute Settings](#)
- [Using CREATE DISKGROUP with Compatibility Attributes](#)
- [Using ALTER DISKGROUP with Compatibility Attributes](#)
- [Viewing Compatibility Attribute Settings](#)
- [Features Enabled By Disk Group Compatibility Attribute Settings](#)
- [Reverting Disk Group Compatibility](#)

You can also set the disk group compatibility settings with Oracle Enterprise Manager, Oracle ASM command-line utility (ASCMD), and Oracle ASM Configuration Assistant (ASMCA). See [Chapter 9, "Administering Oracle ASM with Oracle Enterprise Manager"](#), [Chapter 11, "Oracle ASM Configuration Assistant"](#), and [Chapter 12, "Oracle ASM Command-Line Utility"](#).

Note: Advancing the values for disk group compatibility attributes is an irreversible operation. See ["Reverting Disk Group Compatibility"](#) on page 4-37.

See Also: *Oracle Database SQL Language Reference* for more information about the disk group compatibility SQL statements

Valid Combinations of Compatibility Attribute Settings

When setting the values for the disk group attributes, specify at least the major and minor versions of a valid Oracle Database release number. For example, you can specify compatibility as '11.1' or '11.2'; Oracle assumes that any missing version number digits are zeros.

[Table 4–2](#) shows an example of some of the valid combinations of the disk group compatibility attributes and the valid Oracle ASM and database instance versions for each combination.

Table 4–2 Examples of Disk Group Compatibility Attribute Settings

COMPATIBLE.ASM	COMPATIBLE.RDBMS	COMPATIBLE.ADVM	ASM Instance Version	COMPATIBLE Setting for RDBMS Instance
10.1	10.1	n/a	>= 10.1	>= 10.1
11.1	10.1	n/a	>= 11.1	>= 10.1
11.2	11.1	11.2	>= 11.2	>= 11.1
11.2	11.2	11.2	>= 11.2	>= 11.2

These are some of the possible combinations of Oracle ASM and database releases:

- The database release is 11g release 2 (11.2) (`COMPATIBLE` set to 11.2) and the Oracle ASM release is 11g release 2 (11.2). The `COMPATIBLE.ASM` attribute is set to 10.1 and the Oracle ASM functionality remains at 10g.

- The database release is 10g and the Oracle ASM release is 11.1. Both the `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` disk group attributes are set to the default value of 10.1. The Oracle ASM disk group functionality remains at 10g.
- The database release is 10g and the Oracle ASM release is 11.2. `COMPATIBLE.ASM` is set to 11.2 and `COMPATIBLE.RDBMS` is set to 10.1. The Oracle ASM disk group attributes are displayed in the `V$ASM_ATTRIBUTE` view. Additional disk group functionality is described in [Table 4-3](#) on page 4-36.
- The database release is 11.2 (`COMPATIBLE` set to 11.2) and the Oracle ASM release is 11.2. All the disk group compatibility attributes are set to 11.2. The Oracle ASM disk group attributes are displayed in the `V$ASM_ATTRIBUTE` view. Additional disk group functionality is described in [Table 4-3](#) on page 4-36.

Using CREATE DISKGROUP with Compatibility Attributes

You can specify the compatibility settings for a disk group with the `CREATE DISKGROUP` statement when creating the disk group.

The following example creates a normal redundancy disk group `data1` with the Oracle ASM compatibility set to 11.2 and the RDBMS compatibility set to the default (assuming that the `COMPATIBLE.RDBMS` default is less than or equal to 11.2):

```
CREATE DISKGROUP data1 DISK '/dev/sd*'
  ATTRIBUTE 'compatible.asm' = '11.2';
```

The following example creates a normal redundancy disk group `data2` with the ASM, RDBMS, and ADVM compatibility set to 11.2:

```
CREATE DISKGROUP data2 DISK '/dev/sd*'
  ATTRIBUTE 'compatible.asm' = '11.2', 'compatible.rdbms' = '11.2',
  'compatible.advm' = '11.2';
```

Using ALTER DISKGROUP with Compatibility Attributes

After a disk group has been created, you can use the `ALTER DISKGROUP SQL` statement to change the compatibility attributes. The `ALTER DISKGROUP SQL` statement ensures that Oracle can advance the compatibility of the specified disk group before committing the change.

All of the affected databases and file systems should be online when running `ALTER DISKGROUP` to ensure that advancing compatibility does not reduce the database and file system access. When advancing disk group compatibility, you must advance the `COMPATIBLE.ASM` attribute before the `COMPATIBLE.RDBMS` or `COMPATIBLE.ADVM` attribute to ensure a valid combination of compatible attribute settings as shown in [Table 4-2](#). You can advance only one compatibility attribute in a single `ALTER DISKGROUP` statement.

The following example advances the Oracle ASM compatibility for disk group `data3` to 11.2. An Oracle ASM instance must be at release 11.2 or higher to access the `data3` disk group.

```
ALTER DISKGROUP data3 SET ATTRIBUTE 'compatible.asm' = '11.2';
```

The following example advances the RDBMS and ADVM compatibility of the disk group `data3` to 11.2. This example assumes that the ASM compatibility value is set to 11.2.

```
ALTER DISKGROUP data3 SET ATTRIBUTE 'compatible.rdbms' = '11.2',
  'compatible.advm' = '11.2';
```

Viewing Compatibility Attribute Settings

You can view the disk group compatibility settings in the `V$ASM_ATTRIBUTE` view. However, the `V$ASM_ATTRIBUTE` view does not display any rows when the `COMPATIBLE.ASM` value is set to 10.1. Instead you can determine the values for the `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` disk group compatibility attributes with the `COMPATIBILITY` and `DATABASE_COMPATIBILITY` columns of the `V$ASM_DISKGROUP` view.

See [Example 6-1, "Viewing Disk Group Attributes With V\\$ASM_ATTRIBUTE"](#) on page 6-2 for an example querying the `V$ASM_ATTRIBUTE` view.

You can also display the disk group compatibility attributes with the `ASMCMD` command `lsattr`. For information about `lsattr`, see "[lsattr](#)" on page 12-33.

See Also:

- *Oracle Database SQL Language Reference* for information about disk group attributes
- *Oracle Database Reference* for details on the `V$ASM_DISKGROUP` and `V$ASM_ATTRIBUTE` dynamic performance views

Features Enabled By Disk Group Compatibility Attribute Settings

[Table 4-3](#) describes the features enabled by valid combinations of the disk group compatibility attribute settings.

Table 4-3 Features Enabled by Disk Group Compatibility Attribute Settings

Disk Group Features Enabled	COMPATIBLE.ASM	COMPATIBLE.RDBMS	COMPATIBLE.ADV
Support for larger AU sizes (32 or 64 MB)	>= 11.1	>= 11.1	n/a
Attributes are displayed in the <code>V\$ASM_ATTRIBUTE</code> view	>= 11.1	n/a	n/a
Fast mirror resync	>= 11.1	>= 11.1	n/a
Variable size extents	>= 11.1	>= 11.1	n/a
Exadata storage	>= 11.1.0.7	>= 11.1.0.7	n/a
Intelligent Data Placement	>= 11.2	>= 11.2	n/a
OCR and voting disks in a disk group	>= 11.2	n/a	n/a
Sector size set to non-default value	>= 11.2	>= 11.2	n/a
Oracle ASM SPFILE in a disk group	>= 11.2	n/a	n/a
Oracle ASM File Access Control	>= 11.2	>= 11.2	n/a
Volumes in disk groups	>= 11.2	n/a	>= 11.2

The following list applies to [Table 4-3](#).

- Oracle ASM features not explicitly listed in [Table 4-3](#) do not require advancing the disk group compatibility attribute settings.
- The value of `COMPATIBLE.ASM` must always be greater than or equal to the value of `COMPATIBLE.RDBMS`.
- A value of not applicable (n/a) means that the setting of the attribute has no effect on the feature.

Reverting Disk Group Compatibility

Advancing the values for disk group compatibility attributes is an irreversible operation. If you advance the disk group compatibility settings, you cannot change the values back to the previous settings. To revert to the previous values, you must create a new disk group with the old compatibility attribute settings and then restore the database files that were in the disk group to the new disk group.

When you revert to a new disk group with the old compatibility attribute settings, the latest Oracle ASM features might not be available. For example, if you revert the disk group compatibility to a pre-11.2 value, Oracle ACFS functionality is not available.

For example, you could perform the following procedure to revert a disk group to previous compatibility settings:

1. If the Oracle ASM SPFILE is in the disk group, move this SPFILE out of the disk group:

- a. Connect with SQL*Plus to the Oracle ASM instance.
- b. Create a PFILE in the file system.

For example:

```
SQL> CREATE PFILE '$ORACLE_HOME/dbs/asmospfile.ora' FROM SPFILE
```

2. If the OCR and voting disks are in the disk group, move them out of this disk group.

See Also: The *Oracle Clusterware Administration and Deployment Guide* for information about administering OCR and voting disks

3. Back up any files that you want to save.

- a. Back up the database files.
- b. If an Oracle ACFS file system is mounted on an Oracle ADVM volume on the disk group, the operating system files in the file system must be backed up or copied to a location outside the file system mount point.

4. Create a new disk group using SQL `CREATE DISKGROUP` specifying the previous values for the disk group attribute settings.

For information about creating a disk group, see "[Using the CREATE DISKGROUP SQL Statement](#)" on page 4-3.

5. Restore the database files into the newly created disk group using Recovery Manager (RMAN).

For information about moving data files between disk groups, see "[Moving Data Files Between Oracle ASM Disk Groups Using RMAN](#)" on page 8-9.

6. Drop the disk group with the advanced disk group compatibility settings using SQL `DROP DISKGROUP INCLUDING CONTENTS` to remove the disk group and its contents.

This SQL statement also removes any Oracle ACFS file system and its contents.

For information about dropping a disk group, see "[Dropping Disk Groups](#)" on page 4-44.

Considerations When Setting Disk Group Compatibility in Replicated Environments

If you advance disk group compatibility, then you could enable the creation of files that are too large to be managed by a previous Oracle database release. You must be aware of the file size limits because replicated sites cannot continue using the software from a previous release to manage these large files. The disk group compatibility settings should be the same for all replicated environments.

[Table 4–4](#) shows the maximum Oracle ASM file sizes supported for `COMPATIBLE.RDBMS` settings when the `AU_SIZE` attribute is set to one megabyte for a disk group.

Table 4–4 Maximum Oracle ASM File Sizes for Disk Groups With `AU_SIZE` Equal to 1 MB

Redundancy	<code>COMPATIBLE.RDBMS = 10.1</code>	<code>COMPATIBLE.RDBMS >= 11.1</code>
External	16 TB	140 PB
Normal	5.8 TB	23 PB
High	3.9 TB	15 PB

[Table 4–4](#) shows that Oracle Database 10g can only support a file size of up to 16 TB for external redundancy. If you advance the `COMPATIBLE.RDBMS` attribute to 11.1 or greater, then a file can grow beyond 16 TB. However, the larger size will make the file unusable in a replicated and disaster recovery site if the disaster recovery site has a disk group `COMPATIBLE.RDBMS` setting that is incompatible with the larger size.

For information about Oracle ASM storage sizes, see "[Oracle ASM Storage Limits](#)" on page 4-31.

See Also:

- *Oracle Database Upgrade Guide* for information about database compatibility
- *Oracle Database Administrator's Guide* for information about the `COMPATIBLE` initialization parameter and irreversible compatibility
- *Oracle Database Reference* for information about the `COMPATIBLE` initialization parameter

Managing Oracle ASM File Access Control for Disk Groups

Oracle ASM File Access Control provides optional protection for the content of Oracle ASM disk groups from accidental access by unauthorized Oracle ASM clients, such as an unauthorized database.

To set up Oracle ASM File Access Control, you need to designate separate operating system groups as described in "[Using Separate Operating System Groups for Oracle ASM Users](#)" on page 3-24. Oracle ASM File Access Control is available for Linux and UNIX computers only.

You can manage Oracle ASM file access control with `ASMCMD` commands, Oracle Enterprise Manager, and SQL statements.

This section contains these topics:

- [About Oracle ASM File Access Control](#)

- [Using SQL Statements to Set Disk Group Attributes for Oracle ASM File Access Control](#)
- [Using SQL Statements to Manage Oracle ASM File Access Control](#)

For information about managing Oracle ASM File Access Control with ASMCMD commands, see "[ASMCMD File Access Control Commands](#)" on page 12-50.

For information about managing Oracle ASM File Access Control with Oracle Enterprise Manager, see "[Managing Oracle ASM File Access Control with Oracle Enterprise Manager](#)" on page 9-16.

For information about views that provide details about Oracle ASM file access control, see "[Viewing Oracle ASM File Access Control Information](#)" on page 6-4.

For information about controlling accessing to Oracle ASM instances, see "[Authentication for Accessing Oracle ASM Instances](#)" on page 3-22.

About Oracle ASM File Access Control

Oracle ASM File Access Control restricts the access of files to specific Oracle ASM clients that connect as SYSDBA. An Oracle ASM client is typically a database, which is identified as the user that owns the database instance home. Oracle ASM File Access Control uses this user name to identify a database. Oracle ASM File Access Control restricts access based on the operating system effective user identification number of a database owner. For example, in [Table 3-2, "Separated Operating System Groups and Privileges for Oracle ASM Users"](#) on page 3-25 the databases are identified as oracle1 and oracle2.

Oracle ASM uses file access control to determine the additional privileges that are given to a database that has been authenticated AS SYSDBA on the Oracle ASM instance. These additional privileges include the ability to modify and delete certain files, aliases, and user groups.

You can set up user groups to specify the list of databases that share the same access permissions to Oracle ASM files. User groups are lists of databases and any database that authenticates AS SYSDBA can create a user group. However, only the creator of a group can delete it or modify its membership list.

Each Oracle ASM file has three categories of privileges: owner, group, and other. Each category can have no permission, read-only permission, or read-write permission.

The file owner is usually the creator of the file and has the ability to assign permissions for the file in any of the owner, group, or other categories. The owner can also change the group associated with the file.

When administering Oracle ASM File Access Control, Oracle recommends that you connect as SYSDBA to the database instance that is the owner, or planned owner, of the files in the disk group.

If you want to set up Oracle ASM File Access Control for files in a disk group, perform the following steps:

1. Alter a new or existing disk group to set the Oracle ASM File Access Control disk group attributes.

For a newly-created disk group, you should set the disk group attributes before creating any files in the disk group.

See "[Using SQL Statements to Set Disk Group Attributes for Oracle ASM File Access Control](#)" on page 4-40.

2. For files that exist in a disk group before setting the Oracle ASM File Access Control disk group attributes, you must explicitly set the permissions and ownership on those existing files.

Ensure that the user exists before setting ownership or permissions on a file. The file must be closed before setting the ownership or permissions.

See `ALTER DISKGROUP SET PERMISSION` and `ALTER DISKGROUP SET OWNERSHIP` in "Using SQL Statements to Set Disk Group Attributes for Oracle ASM File Access Control" on page 4-40.

3. Optionally, you can create user groups that are groups of database users that share the same access permissions to Oracle ASM files.

See `ALTER DISKGROUP ADD USERGROUP` in "Using SQL Statements to Set Disk Group Attributes for Oracle ASM File Access Control" on page 4-40.

Using SQL Statements to Set Disk Group Attributes for Oracle ASM File Access Control

To manage Oracle 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 with the `ALTER DISKGROUP` SQL statement.

When you set up file access control on an existing disk group, the files previously created remain accessible by everyone, unless you run the `ALTER DISKGROUP SET PERMISSION` SQL statement to restrict the permissions.

The `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` disk group attributes must be set to 11.2 or higher to enable Oracle ASM File Access Control.

The disk group attributes that control Oracle ASM File Access Control are the following:

- `ACCESS_CONTROL.ENABLED`

This attribute determines whether Oracle ASM File Access Control is enabled for a disk group.

The value can be `true` or `false`. The default is `false`.

If the attribute is set to `true`, accessing Oracle ASM files is subject to access control. If `false`, any user can access every file in the disk group. All other operations behave independently of this attribute.

- `ACCESS_CONTROL.UMASK`

This attribute determines which permissions are masked out on the creation of an Oracle 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.

Setting to 0 masks out nothing. Setting to 2 masks out write permission. Setting to 6 masks out both read and write permissions.

Before setting the `ACCESS_CONTROL.UMASK` disk group attribute, you must set the `ACCESS_CONTROL.ENABLED` attribute to `true` to enable Oracle ASM File Access Control.

Example 4-7 shows how to enable Oracle ASM File Access Control for a disk group with SQL*Plus. In this example, the permissions setting is 026 which enables read-write access for the owner, read access for users in the group, and no access to others not in the group

Example 4-7 Setting Up Oracle ASM File Access Control

```
ALTER DISKGROUP data1 SET ATTRIBUTE 'access_control.enabled' = 'true';
ALTER DISKGROUP data1 SET ATTRIBUTE 'access_control.umask' = '026';
```

Using SQL Statements to Manage Oracle ASM File Access Control

You can use the `ALTER DISKGROUP` SQL statement to manage file access control for Oracle ASM disk groups. These SQL statements are available for both database and Oracle ASM instances.

The SQL statements that support disk group access control are:

- `ALTER DISKGROUP ADD USERGROUP ... WITH MEMBER`

Adds an Oracle ASM user group to a disk group. The user group name is limited to a maximum of 30 characters. The databases identified in the `MEMBER` clause must be in the disk group, as shown by `V$ASM_USER`, or the command returns an error. Any users authenticated as `SYSASM` or `SYSDBA` can create new user groups. For example:

```
SQL> SELECT group_number, os_name FROM V$ASM_USER;
```

```
GROUP_NUMBER OS_NAME
```

```
-----
```

```
1 oracle1
```

```
1 oracle2
```

```
...
```

```
SQL> ALTER DISKGROUP data ADD USERGROUP 'test_grp1'
      WITH MEMBER 'oracle1','oracle2';
```

- `ALTER DISKGROUP DROP USERGROUP`

Drops an Oracle ASM user group from a disk group. Dropping a group might leave some files without a valid group. If you want those files to have a valid group, you must manually update the group associated with those files to a valid group.

```
SQL> ALTER DISKGROUP data DROP USERGROUP 'test_grp1';
```

- `ALTER DISKGROUP MODIFY USERGROUP ADD MEMBER`

Adds users to the specified user group. The users must be in the disk group, as shown by `V$ASM_USER`, or the command returns an error. Only the creator of the group or the Oracle ASM administrator can modify group membership.

```
SQL> ALTER DISKGROUP data MODIFY USERGROUP 'test_grp2' ADD MEMBER 'oracle2';
```

- `ALTER DISKGROUP MODIFY USERGROUP DROP MEMBER`

Removes users from the specified user group. If a member is not in the user group, then an error is returned. Only the creator of the group or the Oracle ASM administrator can modify group membership.

```
SQL> ALTER DISKGROUP data MODIFY USERGROUP 'test_grp2' DROP MEMBER 'oracle2';
```

- `ALTER DISKGROUP ADD USER`

Adds operating system (OS) users to an Oracle ASM disk group, so that these users can have access privileges on the disk group. The users must be existing operating system users, and their user names must have a corresponding operating system user ID or system ID. If a user exists in the disk group, as shown

by `V$ASM_USER`, then the command records an error and continues to add other users, if any.

The operating system user of a running database instance is automatically added to a disk group when the database instance accesses that disk group and creates files. However, if you want a database instance to read files in a disk group without creating any files, then you must use this SQL statement to add that database user to the disk group. Also, you can use this statement to add a database user to an existing disk group immediately after setting the Oracle ASM File Access Control disk group attributes and before creating new files.

```
SQL> ALTER DISKGROUP DATA ADD USER 'oracle1';
```

- ALTER DISKGROUP DROP USER

Drops operating system users from an Oracle ASM disk group. If a user is not in the disk group, then this command records an error and continues to drop other users, if any.

If the user owns any files on the same Oracle ASM disk group, then this command fails with an error, unless the `CASCADE` keyword is specified. If the latter case, then the user is deleted, along with all the files that the user owns.

If any files owned by the user are currently open, then the `DROP USER` command fails, and no files are deleted.

```
SQL> ALTER DISKGROUP DATA DROP USER 'oracle1';
```

- ALTER DISKGROUP SET PERMISSION

Modifies permissions of an Oracle ASM file. Setting `read only` permission to a file that has `read write` permission revokes the `write` permission. Only the file owner or the Oracle ASM administrator can change the permissions of a file. You cannot change the permissions on an open file.

```
SQL> ALTER DISKGROUP data SET PERMISSION OWNER=read write, GROUP=read only,
      OTHER=none FOR FILE '+data/controlfile.f';
```

- ALTER DISKGROUP SET OWNERSHIP

Changes the owner or group of a file to the specified user or user group name, respectively. If the specified user or user group name does not exist, this command fails with an error. Only the owner of the file or the Oracle ASM administrator can issue this command, and only the Oracle ASM administrator can change the owner. Also, the user group name must exist, and the owner of the file must be a member of that group. You cannot change the ownership of an open file.

```
SQL> ALTER DISKGROUP data SET OWNERSHIP OWNER='oracle1', GROUP='test_grp1'
      FOR FILE '+data/controlfile.f';
```

See Also: The `ALTER DISKGROUP` SQL statement in the *Oracle Database SQL Language Reference*

Mounting and Dismounting Disk Groups

Disk groups that are specified in the `ASM_DISKGROUPS` initialization parameter are mounted automatically at Oracle ASM instance startup. This makes them available to all database instances running on the same node as Oracle ASM. The disk groups are dismounted at Oracle ASM instance shutdown. Oracle ASM also automatically

mounts a disk group when you initially create it, and dismounts a disk group if you drop it.

When a disk group is mounted, a disk group number is chosen. This number may change across disk group mounts. A disk group number is not recorded in any persistent structure, but the current value can be viewed in the `GROUP_NUMBER` column of the `V$ASM` views.

There might be times that you want to mount or dismount disk groups manually. For these actions use the `ALTER DISKGROUP . . . MOUNT` or `ALTER DISKGROUP . . . DISMOUNT` statement. You can mount or dismount disk groups by name, or specify `ALL`.

If you try to dismount a disk group that contains open files, the statement fails, unless you also specify the `FORCE` clause.

In a clustered Oracle ASM environment in `RESTRICTED` mode, a disk group is mounted in single-instance exclusive mode. No other Oracle ASM instance in that cluster can mount that disk group. In this mode the disk group is not usable by any Oracle ASM client. Use this mode to perform a fast rebalance.

The following SQL statement dismounts all disk groups that are currently mounted to the Oracle ASM instance:

```
ALTER DISKGROUP ALL DISMOUNT;
```

The following SQL statement mounts disk group `data1`:

```
ALTER DISKGROUP data1 MOUNT;
```

Mounting Disk Groups Using the FORCE Option

Oracle ASM provides a `MOUNT FORCE` option with `ALTER DISKGROUP` to enable Oracle ASM disk groups to be mounted in normal or high redundancy modes even though some Oracle ASM disks may be unavailable to the disk group at mount time. The default behavior without the `FORCE` option is to fail to mount a disk group that has damaged or missing disks.

The `MOUNT FORCE` option is useful in situations where a disk is temporarily unavailable and you want to mount the disk group with reduced redundancy while you correct the situation that caused the outage.

To successfully mount with the `MOUNT FORCE` option, Oracle ASM must be able to find at least one copy of the extents for all of the files in the disk group. In this case, Oracle ASM can successfully mount the disk group, but with potentially reduced redundancy.

The disks that Oracle ASM is unable to access are put in an offline mode. Oracle ASM then begins timing the period that these disks are in an offline mode. If the disk offline time period exceeds the timer threshold set by `DISK_REPAIR_TIME` disk group attribute, then those disks are permanently dropped from the disk group. You can change the offline timer after a disk is put in an offline state by using the `ALTER DISKGROUP OFFLINE` statement. For more information about setting the `DISK_REPAIR_TIME` disk group attribute, see ["Oracle ASM Fast Mirror Resync"](#) on page 4-27.

Note: An Oracle ASM instance mounts an incomplete disk group differently depending on the specified compatibility. See ["Features Enabled By Disk Group Compatibility Attribute Settings"](#) on page 4-36.

In clustered Oracle ASM environments, if an Oracle ASM instance is not the first instance to mount the disk group, then using the `MOUNT FORCE` statement fails. This is because the disks have been accessed by another instance and the disks are not locally accessible.

If all disks are available, then using the `FORCE` option causes the `MOUNT` command to fail. This discourages unnecessary and improper use of the feature.

The following example shows how to use the `FORCE` option to force the mount of the `data1` disk group:

```
ALTER DISKGROUP data1 MOUNT FORCE
```

See Also: The *Oracle Database SQL Language Reference* for additional information about the `ALTER DISKGROUP` statement and the `FORCE` option

Checking the Internal Consistency of Disk Group Metadata

You can check the internal consistency of disk group metadata using the `ALTER DISKGROUP` statement with the `CHECK` keyword. You can use this statement to check specific files in a disk group, specific disks or all disks in a disk group, or specific failure groups within a disk group. The disk group must be mounted to perform these checks.

By default, the `CHECK DISK GROUP` clause verifies all of the metadata directories. Oracle ASM displays summary errors and writes the details about the errors in an alert log. The `CHECK` keyword performs the following operations:

- Verifies the consistency of the disk
- Cross checks all of the file extent maps and allocation tables for consistency
- Checks that the alias metadata directory and file directory are linked correctly
- Verifies that the alias directory tree is linked correctly
- Checks that Oracle ASM metadata directories do not have unreachable allocated blocks

The `REPAIR | NOREPAIR` clause specifies whether Oracle ASM should attempt to repair errors that are found during the check. The default is `NOREPAIR`. Use the `NOREPAIR` clause to receive alerts about inconsistencies and to suppress Oracle ASM from resolving the errors automatically. The following example statement checks for consistency in the metadata for all disks in the `data1` disk group:

```
ALTER DISKGROUP data1 CHECK ALL;
```

See Also: The *Oracle Database SQL Language Reference* for additional information about the `CHECK` clause syntax

Dropping Disk Groups

The `DROP DISKGROUP` statement enables you to delete an Oracle ASM disk group and optionally, all of its files. You can specify the `INCLUDING CONTENTS` clause if you also want to delete all files that are contained in the disk group. The default is `EXCLUDING CONTENTS`, which provides syntactic consistency and prevents you from dropping the disk group if it has any contents

The Oracle ASM instance must be started and the disk group must be mounted with none of the disk group files open, in order for the `DROP DISKGROUP` statement to succeed. The statement does not return until the disk group has been dropped.

When you drop a disk group, Oracle ASM dismounts the disk group and removes the disk group name from the `ASM_DISKGROUPS` initialization parameter if a server parameter file is being used. If a text initialization parameter file is being used, and the disk group is mentioned in the `ASM_DISKGROUPS` initialization parameter, then you must remove the disk group name from the `ASM_DISKGROUPS` initialization parameter before the next time that you shut down and restart the Oracle ASM instance.

The following statement deletes `data1`:

```
DROP DISKGROUP data1;
```

After ensuring that none of the files contained in `data1` are open, Oracle ASM rewrites the header of each disk in the disk group to remove Oracle ASM formatting information. The statement does not specify `INCLUDING CONTENTS`, so the drop operation fails if the disk group contains any files.

If an Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file system is mounted on a volume contained in the disk group, then the file system must be dismounted. If the file system has been registered, then it must be deregistered. The `INCLUDING CONTENTS` clause must be used to drop this disk group. All data in the file system is destroyed. To view the volumes and mount paths associated with a disk group, you can query the `V$ASM_VOLUME` view. For an example of a query on the `V$ASM_VOLUME` view, see [Example 6-14](#) on page 6-7. For information about deregistering and dismounting Oracle ACFS file systems, see "[Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#)" on page 13-5.

If you cannot mount a disk group but need to drop it, you can use the `FORCE` option of the `DROP DISKGROUP` statement. This command enables you to remove the headers on disks that belong to a disk group that cannot be mounted by any Oracle ASM instances as shown in the following example:

```
DROP DISKGROUP data1 FORCE
```

The disk group on which you perform this operation should not be mounted anywhere in the cluster. When you use the `FORCE` option, the Oracle ASM instance does not attempt to verify that a disk group is being used by another Oracle ASM instance in the same storage subsystem.

Note: Use the `FORCE` option with extreme caution.

You can also drop a disk group with Oracle Enterprise Manager. See "[Dropping Disk Groups](#)" on page 9-12.

Renaming Disks Groups

The `renamedg` tool enables you to change the name of a cloned disk group. The disk group must be dismounted on all nodes in the cluster before running `renamedg` on the disk group.

`renamedg` renames a disk group using a two-step process:

1. Phase one

This phase generates a configuration file to be used in phase two.

2. Phase two

This phase uses the configuration file to perform the renaming of the disk group.

The syntax is:

```
renamedg {-help | help=true}
```

```
renamedg
    [phase={one|two|both}] dgroup=diskgroup
    newdgroup=newdiskgroup [config=configfile]
    [asm_diskstring=discoverystring, discoverystring...]
    [clean={true|false}] [check={true|false}]
    [confirm={true|false}] [verbose={true|false}]
    [keep_voting_files={true|false}]
```

- `phase={one|two|both}`

Specifies the phase to be executed. Allowed values are `one`, `two`, or `both`. This argument is optional. The default is `both`.

Typically you would run both phases. If a problem occurs during the second phase, then you can re-run phase `two` using the generated configuration file.
- `dgroup=diskgroup`

Specifies the name of the disk group that to be renamed.
- `newdgroup=newdiskgroup`

Specifies the new name for the disk group.
- `config=configfile`

Specifies the path to the configuration file to be generated during phase one or specifies the path to the configuration file to be used during phase two.

This argument is optional. The default configuration file is named `renamedg_config` and is located in the directory in which the command is executed. The single quotations may be required on some platforms.
- `asm_diskstring=discoverystring, discoverystring...`

Specifies the Oracle ASM discovery strings. The `asm_diskstring` value must be specified if the Oracle ASM disks are not in the default location for the platform. The single quotations may be required on some platforms, usually when wildcard characters are specified.
- `clean={true|false}`

Specifies whether to tolerate errors that are otherwise ignored. The default is `true`.
- `check={true|false}`

Specifies a boolean value that is used in the second phase. If `true`, then the tool prints the list of changes that will be made to the disks. No writes are issued. It is an optional parameter that defaults to `false`.
- `confirm={true|false}`

Specifies a boolean value that is used in the second phase. If `false`, then the tool prints the changes that will be made and seeks confirmation before actually

making the changes. It is an optional value that defaults to `false`. If `check` is set to `true`, then the value of this parameter is redundant.

- `verbose={true|false}`

Specifies verbose execution when `verbose=true`. The default is `false`.

- `keep_voting_files={true|false}`

Specifies whether voting files are kept in the renamed disk group. The default is `false` which deletes the voting files from the renamed disk group.

Note: `renamedg` does not update resources, nor does `renamedg` update any file references within the database. Because of this, the original disk group resource is not automatically deleted after the completion of phase two. The status of the old disk group resource can be checked with the Oracle Clusterware Control (CRSCTL) `crsctl stat res -t` command and then manually deleted with the Server Control Utility (SRVCTL) `svrctl remove diskgroup` command.

[Example 4-8](#) shows several examples of the use of `renamedg`. The first example renames the data disk group to `new_data` using a disk string to locate the disks and the `verbose` option is enabled. The second example only creates a configuration file during the completion of phase one of the `renamedg` operation. The third example runs phase two of the `renamedg` operation using a configuration file generated from a phase one execution of `renamedg`.

Example 4-8 Using renamedg

```
$ renamedg dgname=data newdgname=new_data asm_diskstring='/devices/disk*'
    verbose=true

$ renamedg phase=one dgname=data newdgname=new_data
    asm_diskstring='/devices/disk*' config=/tmp/data2.conf verbose=true

$ renamedg phase=two dgname=data newdgname=new_data config=/tmp/data2.conf
    verbose=true
```

Introduction to Oracle ACFS

This chapter describes Oracle Automatic Storage Management Cluster File System (Oracle ACFS) concepts and provides an overview of Oracle ACFS features.

This chapter contains the following topics:

- [Overview of Oracle ACFS](#)
- [Understanding Oracle ACFS Concepts](#)
- [Understanding Oracle ACFS Administration](#)
- [Overview of Oracle ASM Dynamic Volume Manager](#)

See Also:

- ["Views Containing Oracle ACFS Information"](#) on page 6-6 for information about using views to obtain Oracle ACFS information
- [Chapter 10, "Administering Oracle ACFS With Oracle Enterprise Manager"](#) for information about using Oracle Enterprise Manager to administer Oracle ACFS
- [Chapter 11, "Oracle ASM Configuration Assistant"](#) for information about using Oracle ASM Configuration Assistant to administer Oracle ACFS
- ["ASMCMD Volume Management Commands"](#) on page 12-58 for information about volume management commands
- [Chapter 13, "Oracle ACFS Command-Line Tools"](#) for information about Oracle ACFS utilities

Overview of Oracle ACFS

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a multi-platform, scalable file system, and storage management technology that extends Oracle Automatic Storage Management (Oracle ASM) functionality to support customer files maintained outside of the Oracle Database. Oracle ACFS supports many database and application files, including executables, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

Notes:

- Oracle ASM is the preferred storage manager for all database files. It has been specifically designed and optimized to provide the best performance for database file types.
- Oracle ACFS is the preferred file manager for non-database files. It is optimized for general purpose files.
- Oracle ACFS does not support any file that can be directly stored in Oracle ASM.

Not supported means Oracle Support Services will not take calls and development will not fix bugs associated with storing unsupported file types in Oracle ACFS.

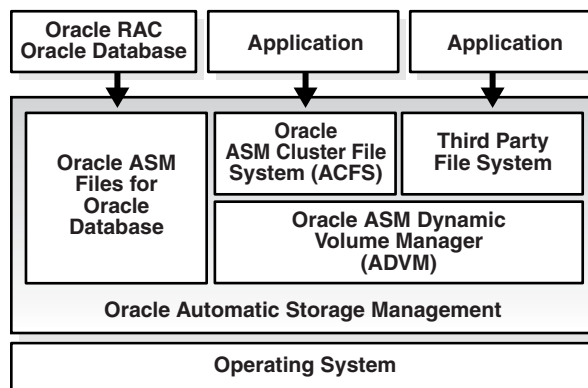
For a list of file types supported by Oracle ASM, see [Table 7-1, "File Types Supported by Automatic Storage Management"](#) on page 7-1.

- Oracle ACFS does not support files for the Oracle grid infrastructure home.

An Oracle ACFS file system layers on Oracle ASM and is configured with Oracle ASM storage, as shown in [Figure 5-1](#). Oracle ACFS leverages Oracle ASM functionality that enables:

- Oracle ACFS dynamic file system resizing
- Maximized performance through direct access to Oracle ASM disk group storage
- Balanced distribution of Oracle ACFS across Oracle ASM disk group storage for increased I/O parallelism
- Data reliability through Oracle ASM mirroring protection mechanisms

Figure 5-1 Oracle Automatic Storage Management Storage Layers



Oracle ACFS establishes and maintains communication with the Oracle ASM instance to participate in Oracle ASM state transitions including Oracle ASM instance and disk group status updates and disk group rebalancing. Oracle Automatic Storage Management with Oracle ACFS and Oracle ASM Dynamic Volume Manager (Oracle ADVM) delivers support for all customer data and presents a common set of Oracle storage management tools and services across multiple vendor platforms and operating system environments on both Oracle Restart (single-node) and cluster configurations. For an overview of Oracle ADVM, see ["Overview of Oracle ASM Dynamic Volume Manager"](#) on page 5-10.

Oracle ACFS is tightly coupled with Oracle Clusterware technology, participating directly in Clusterware cluster membership state transitions and in Oracle Clusterware resource-based high availability (HA) management. In addition, Oracle installation, configuration, verification, and management tools have been updated to support Oracle ACFS.

Oracle ACFS can be accessed and managed using native operating system file system tools and standard application programming interfaces (APIs). Oracle ACFS can be configured and managed using Oracle Enterprise Management tools. Oracle ACFS can be accessed using industry standard Network Attached Storage (NAS) File Access Protocols: Network File System (NFS) and Common Internet File System (CIFS).

In addition to sharing file data, Oracle ACFS provides additional storage management services including support for the Oracle Restart mount registry and the Oracle grid infrastructure clusterwide mount registry, dynamic on-line file system resizing, and multiple space-efficient snapshots for each file system. For information about the mount registry, see "[About the Oracle ACFS Mount Registry](#)" on page 5-6.

Oracle ACFS contributes to the overall Oracle storage management by providing:

- A general-purpose single-node (standalone) and cluster file system solution that is integrated with Oracle ASM, Oracle Clusterware, and Oracle Enterprise Management technologies
- A common set of file system features across multiple vendor platforms and operating systems, offering an alternative to native operating system or third-party file system solutions
- Single-node and clusterwide shared Oracle Database homes and application data
- Uniform, coherent shared file access and clusterwide naming of all customer application files maintained outside of the database

Oracle ACFS accommodates large storage capacities and large numbers of cluster nodes. It efficiently manages large numbers of file systems, files, and supports both small and large sized files with exabyte-capable file and file system capacities. Oracle ACFS provides optimized fast directory lookup for large directories with hundreds of thousands of files.

Oracle ACFS file systems are generally mounted on all CSS cluster members. In the event of a member failure, another cluster member quickly recovers any outstanding metadata transactions on behalf of the failed member. Following recovery, access by other active cluster members and any remote client systems can resume.

Create Oracle Database data files directly on Oracle ASM storage for best performance because Oracle ACFS does not support `directIO` for file read and write operations in 11g Release 2 (11.2).

See [Chapter 6, "Using Views to Display Information"](#) and [Chapter 13, "Oracle ACFS Command-Line Tools"](#).

Understanding Oracle ACFS Concepts

This section describes concepts for the key Oracle ACFS components and contains the following topics:

- [About Oracle ACFS](#)
- [About Oracle ACFS and Oracle Database Homes](#)
- [About Oracle ASM Dynamic Volume Manager](#)

- [About the Oracle ACFS Driver Model](#)
- [About the Oracle ACFS Mount Model and Namespace](#)
- [About the Oracle ACFS Mount Registry](#)
- [About Oracle ACFS Snapshots](#)
- [About Oracle ACFS and Backup and Restore](#)
- [About Oracle ACFS Integration with Oracle ASM](#)

About Oracle ACFS

Oracle ACFS is designed as a general-purpose, single-node and clusterwide file system that delivers support for customer file system application data maintained outside of the Oracle Database. Users and applications can access and manage Oracle ACFS using native operating system file system application programming interfaces (APIs) and command line interface (CLI) tools. Users can also manage Oracle ACFS with the Oracle Enterprise Management and Configuration tools.

Oracle ACFS supports large files with 64-bit file and file system data structure sizes leading to exabyte capable file and file system capacities on 64 bit platforms. Variable extent-based storage allocation, high-performance directories contribute to fast performance, and shared disk configurations that provide direct storage paths to Oracle ACFS file data from each cluster member. File system integrity and fast recovery is achieved through the use of Oracle ACFS metadata checksums and journaling. Oracle ACFS is designed as a multi-node, shared file system model that delivers coherent, cached, direct storage paths to Oracle ACFS file data from each cluster member.

Oracle ACFS files systems are typically configured for clusterwide access. File systems, files, and directories are visible and accessible from all cluster members and can be referenced by users and applications using the same path names from any cluster member. This design enables simplified application deployments across cluster members and facilitates both multiple instance cluster applications and high availability (HA) failover of unmodified single-node applications.

Oracle ACFS presents single system file access semantics across cluster configurations. Applications and users on all cluster members are presented with the same view of shared Oracle ACFS file data at all times supported by the Oracle ACFS clusterwide user and metadata cache coherency mechanism.

About Oracle ACFS and Oracle Database Homes

An Oracle ACFS file system can be used for an Oracle Database home file system. The file system can be configured as any Oracle Database home, as well as a shared or non-shared Oracle Database home in Oracle RAC cluster configurations.

Note: Oracle Database data files are not supported on Oracle ACFS. These data files should be placed in Oracle ASM disk groups.

Oracle Database homes created on Oracle ACFS should be located directly under the `ORACLE_BASE/acfsmounts` mount point location where `ORACLE_BASE` is the user-specified database directory of the database owner. The Oracle Database base (`ORACLE_BASE` for database) directory should not be the Oracle grid infrastructure base (`ORACLE_BASE` for grid) directory or should not be located under the Oracle grid infrastructure base directory (`ORACLE_BASE` for grid).

One or more Oracle Database homes on Oracle ACFS can be created under the mount point. Each Oracle Database home should be created using a separate Oracle ACFS file system located under the `acfsmounts` mount point. For example, `ORACLE_BASE/acfsmounts/dbhome1` and `ORACLE_BASE/acfsmounts/dbhome2` can be configured as two individual file systems mounted on `dbhome1` and `dbhome2` mount points created under the `ORACLE_BASE/acfsmounts` mount point.

Note: In an Oracle grid infrastructure Clusterware configuration, run `srvctl add filesystem` to automount a file system when an Oracle Database home is installed on the Oracle ACFS file system. That file system should not be added to the registry. For information about Server Control Utility (SRVCTL), see *Oracle Real Application Clusters Administration and Deployment Guide*.

After the installation of Grid Infrastructure Software and before the installation of the Oracle Database software with Oracle Universal Installer (OUI), you can create an Oracle ACFS file system to be configured for use as an Oracle Database Home. You can use the Oracle ASM Configuration Assistant (ASMCA) to create the Oracle file system or you can perform the general procedure for creating the file system. See "[Creating an Oracle ACFS File System for a Database Home](#)" on page 11-13 and "[Basic Steps to Manage an Oracle ACFS](#)" on page 13-1.

After the Oracle ACFS file system is created, the Oracle ACFS-based database home mount point location can be selected as the Oracle Database Home location by browsing to and then choosing the directory during the OUI Database Software installation.

Oracle ACFS file systems can be also configured for use as application homes and Oracle Database homes. However, Oracle ACFS file systems cannot be used for an Oracle base directory or an Oracle grid infrastructure home that contains the software for Oracle Clusterware, Oracle ASM, Oracle ACFS, and Oracle ADVM components.

See Also:

- The *Oracle Database Installation Guide* for your environment for information about the installation of Oracle Database software and storage options
- The *Oracle Grid Infrastructure Installation Guide* for your environment for information about the installation of Oracle Database software and storage options
- *Oracle Database Administrator's Guide* for information about using Oracle-managed files

About Oracle ASM Dynamic Volume Manager

The Oracle ASM Dynamic Volume Manager (Oracle ADVM) provides volume management services and a standard disk device driver interface to clients. File systems and other disk-based applications issue I/O requests to Oracle ADVM volume devices as they would to other storage devices on a vendor operating system.

For more information about Oracle ADVM, see "[Overview of Oracle ASM Dynamic Volume Manager](#)" on page 5-10.

About the Oracle ACFS Driver Model

An Oracle ACFS file system is installed as a dynamically loadable vendor operating system (OS) file system driver and tool set that is developed for each supported operating system platform. The driver is implemented as a Virtual File System (VFS) and processes all file and directory operations directed to a specific file system.

Note: Errors encountered by the drivers are written to the native operating system console and system event loggers. See ["Understanding Oracle ACFS I/O Failure Console Messages"](#) on page B-7.

About the Oracle ACFS Mount Model and Namespace

Oracle ACFS is designed as a hierarchical file system containing files and subdirectories organized into a tree-structured namespace with files at the leaf nodes. The namespace design is a single-file system naming model for both single-node and cluster configurations. This enables each cluster member to present shared files to cluster applications using the same path names, simplifying multi-node application and user access, as well as overall file system administration. The Oracle ACFS mount model also accommodates node local mounts and cluster node subset mounts in cluster configurations to accommodate additional customer requirements.

With a primary focus upon the support of customer application files, Oracle ACFS is not for use as root file system of an operating system or boot storage. Otherwise, an Oracle ACFS file system may be mounted into the native operating system file system namespace using either the mount command line tool or an Oracle Enterprise Management tool.

About the Oracle ACFS Mount Registry

The Oracle ACFS mount registry is an operating system independent, multi-platform mount registration facility supporting both Oracle Restart and Oracle grid infrastructure cluster configurations. File systems that are to be mounted persistently (across reboots) can be registered with the Oracle ACFS mount registry. In cluster configurations, registered Oracle ACFS file systems are automatically mounted by the mount registry, similar to a clusterwide mount table. However, automatic mounting of registered Oracle ACFS file systems are not supported for Oracle Restart configurations. For more information, see ["Oracle ACFS and Oracle Restart"](#) on page B-6.

By default, an Oracle ACFS file system that is inserted into the cluster mount registry is automatically mounted on all cluster members, including cluster members that are added subsequent to the registry addition. However, the cluster mount registry also accommodates single-node and multi-node (subset of cluster nodes) file system registrations. The mount registry actions for each cluster member mount only registered file systems that have been designated for mounting on that member.

The Oracle ACFS registry resource actions are designed to automatically mount a file system only once for each Oracle grid infrastructure initialization to avoid potential conflicts with administrative actions to dismount a given file system.

About Oracle ACFS Snapshots

An Oracle ACFS snapshot is an online, read-only, point in time copy of an Oracle ACFS file system. The snapshot copy is space-efficient and uses Copy-On-Write

functionality. Before an Oracle ACFS file extent is modified or deleted, its current value is copied to the snapshot to maintain the point-in-time view of the file system.

Oracle ACFS snapshots are immediately available for use after they are created. They are always online while the file system is mounted. As a result, an Oracle ACFS snapshot can support the online recovery of files inadvertently modified or deleted from a file system. With up to 63 snapshot views supported for each file system, flexible online file recovery solutions spanning multiple views can be employed. An Oracle ACFS snapshot can also be used as the source of a file system backup, as it can be created on demand to deliver a current, consistent, online view of an active file system.

Oracle ACFS snapshot storage is maintained within the file system, eliminating the need to manage separate storage pools for file systems and snapshots. Oracle ACFS file systems can be dynamically resized to accommodate additional file and snapshot storage requirements.

Oracle ACFS snapshots are administered with the `acfsutil snap` commands. For information about the `acfsutil snap` commands, see ["acfsutil snap create"](#) on page 13-30 and ["acfsutil snap delete"](#) on page 13-32.

You can also manage Oracle ACFS snapshots with Oracle Enterprise Manager. For information about using Oracle Enterprise Manager, see ["Managing Oracle ACFS Snapshots with Oracle Enterprise Manager"](#) on page 10-5.

About Oracle ACFS and Backup and Restore

Oracle ACFS executes on operating system platforms as a native file system technology supporting native operating system file system application programming interfaces (APIs). As a result, backup applications that access files using the native operating system file system interfaces are able to access and backup Oracle ACFS file systems as well as other native operating system file systems. Oracle ACFS snapshots can be dynamically created and used to present a consistent, on-line view of an active file system to a backup application.

Backup applications that use interfaces other than the standard operating system interfaces (read or write) are not supported with Oracle ACFS. For example, Windows backup applications that depend upon the presence of reparse points or the Windows Volume Shadow Copy Service (VSS) are not supported.

About Oracle ACFS Integration with Oracle ASM

Oracle ACFS is always configured with Oracle ASM storage and interfaces with Oracle ASM storage through a traditional device file. This device file is presented by the Oracle ASM Dynamic Volume Manager (Oracle ADVM) and is constructed using a dynamic volume file. The Oracle ADVM volume device file is created automatically following the creation of an Oracle ADVM volume. An Oracle ACFS file system is then bound to the Oracle ADVM device file during the file system creation. After an Oracle ACFS is configured and mounted, the file system inherits the Oracle ASM storage management features associated with an Oracle ADVM volume, including dynamic balanced distribution, mirroring and striping, and dynamic resizing.

The Oracle ACFS driver establishes communication with the Oracle ASM instance to receive Oracle ASM status information including Oracle ASM instance and disk group state transitions.

For information about Oracle ACFS and Oracle ASM operations, see ["Oracle ACFS and Dismount or Shutdown Operations"](#) on page 5-10.

Understanding Oracle ACFS Administration

This section describes Oracle ACFS administration and contains the following topics:

- [Oracle ACFS and File Access and Administration Security](#)
- [Oracle ACFS and Grid Infrastructure Installation](#)
- [Oracle ACFS and Grid Infrastructure Configuration](#)
- [Clusterware Resources and Oracle ACFS Administration](#)
- [Oracle ACFS and Dismount or Shutdown Operations](#)

Oracle ACFS and File Access and Administration Security

Oracle ACFS supports both traditional UNIX-style file access control classes (user, group, other) for Linux and UNIX platforms and the Windows Security Model including file access control lists (ACLs) for Windows platforms. Most Oracle ACFS administrative actions are performed by users with either root or Oracle ASM administration privileges for Linux and UNIX platforms and by users with Windows Administrative privileges on Windows platforms. General Oracle ACFS information for file systems can be accessed by any system user.

In support of Oracle ACFS administration, Oracle recommends that the Oracle ASM administrator role is given to a root privileged user, as many common Oracle ACFS file system management tasks including mount, umount, fsck, driver load, and driver unload are root privileged operations. Other privileged Oracle ACFS file system operations that do not require root privileges can be performed by the Oracle ASM administrator. If the Oracle ASM administrator role is not given to a root privileged user, access to Oracle ACFS file systems can be restricted with the `norootsuid` and `nodev` mount options.

For information about Oracle ASM privileges, see "[About Privileges for Oracle ASM](#)" on page 3-23. For information about administering Oracle ACFS, see [Chapter 10, "Administering Oracle ACFS With Oracle Enterprise Manager"](#) and [Chapter 13, "Oracle ACFS Command-Line Tools"](#).

Oracle ACFS and Grid Infrastructure Installation

Oracle grid infrastructure includes Oracle Clusterware, Oracle ASM, Oracle ACFS, Oracle ADVM, and driver resources software components, which are installed into the Grid Infrastructure Home using the Oracle Universal Installation (OUI) tool.

For information about Oracle ACFS and Oracle Clusterware resources, see "[Clusterware Resources and Oracle ACFS Administration](#)" on page 5-9.

Oracle ACFS and Grid Infrastructure Configuration

After a Grid Infrastructure installation, Oracle Clusterware is operational and the Oracle ASM Configuration Assistant (ASMCA) can be used to start the Oracle ASM instance and create Oracle ASM disk groups, Oracle ADVM volumes, and Oracle ACFS file systems. Alternatively, Oracle ASM disk groups and Oracle ADVM volumes can be created using SQL*Plus, ASMCMD command line tools, or Oracle Enterprise Manager. File systems can be created using operating system command-line tools or Oracle Enterprise Manager.

Oracle ACFS file systems are configured with Oracle ADVM based operating system storage devices that are created automatically following the creation of an Oracle ADVM dynamic volume file. After a volume file and its associated volume device file

are created, a file system can be created and bound to that operating system storage device. Following creation, an Oracle ACFS file system can be mounted, after which it is accessible to authorized users and applications executing file and file system operations.

For an example of the specific actions required to create a new file system, see ["Basic Steps to Manage an Oracle ACFS"](#) on page 13-1. For information about managing Oracle ACFS file systems with ASMCA, see ["Managing Oracle ACFS File Systems With Oracle ASM Configuration Assistant"](#) on page 11-11. For information about managing Oracle ACFS file systems with Oracle Enterprise Manager, see ["Creating Oracle ACFS Volumes and File Systems"](#) on page 10-1.

Clusterware Resources and Oracle ACFS Administration

Oracle Clusterware resources support Oracle ACFS, Oracle Kernel Services Driver (OKS), Oracle ADVN startup, the Oracle ACFS cluster mount registry, and Oracle ACFS single file system startup, shutdown, and steady-state actions.

The following list summarizes Oracle ACFS resource-based management.

- The Oracle ACFS, Oracle Kernel Services (OKS), and Oracle ADVN drivers are dynamically loaded when the Oracle ASM instance is started.
 - Oracle ACFS
 - This driver processes all Oracle ACFS file and directory operations.
 - Oracle ADVN
 - This driver provides block device services for Oracle ASM volume files that are used by file systems for creating file systems.
 - Oracle Kernel Services Driver (OKS)
 - This driver provides portable driver services for memory allocation, synchronization primitives, and distributed locking services to Oracle ACFS and Oracle ADVN.

The drivers are managed as a single resource set.

For information, see ["Oracle ACFS Drivers Resource Management"](#) on page B-3.

- Oracle ACFS file systems listed in the Oracle ACFS mount registry are automatically mounted during Grid Infrastructure initialization and as new mount registry entries are created.

The registry resource is used to manage activation of the Oracle ACFS mount registry and to support the mount and dismount actions for Oracle ACFS file systems listed in the Oracle ACFS mount registry.

For information, see ["Oracle ACFS Registry Resource Management"](#) on page B-4.

- Individual file systems are either manually mounted or dismounted using an Oracle ACFS or Oracle Clusterware command-line tool, or automatically mounted or dismounted as a result of a resource dependency action.

For example, a file system hosting an Oracle Database home can be named in the dependency list of the associated Oracle Database resource such that issuing a start on the database resource results in mounting the dependent Oracle ACFS hosted database home file system.

For information, see ["Oracle ACFS Individual File System Resource Management"](#) on page B-6.

Oracle ACFS and Dismount or Shutdown Operations

It is important to dismount any active file system configured with an Oracle ADVM volume device file before an Oracle ASM instance is shutdown or a disk group is dismounted. After the file systems are dismounted, all open references to Oracle ASM files are removed and associated disk groups can be dismounted or the instance shut down.

If the Oracle ASM instance or disk group is forcibly shut down or fails while an associated Oracle ACFS is active, the file system is placed into an offline error state. If any file systems are currently mounted on Oracle ADVM volume files, the `SHUTDOWN ABORT` command should not be used to terminate the Oracle ASM instance without first dismounting those file systems. Otherwise, applications encounter I/O errors and Oracle ACFS user data and metadata being written at the time of the termination may not be flushed to storage before the Oracle ASM storage is fenced. If it is not possible to dismount the file system, then you should issue two `sync (1)` commands to flush cached file system data and metadata to persistent storage before issuing the `SHUTDOWN ABORT` operation.

Any subsequent attempt to access an offline file system returns an error. Recovering a file system from that state requires dismounting and remounting the Oracle ACFS file system. Dismounting an active file system, even one that is offline, requires stopping all applications using the file system, including any shell references. For example, a previous change directory (`cd`) into a file system directory. The Linux `fuser` or `lsof` commands or Windows `handle` command list information about processes and open files.

For information about shutting down an Oracle ASM instance, see ["Shutting Down an Oracle ASM Instance"](#) on page 3-16. For information about dismounting a disk group, see ["Mounting and Dismounting Disk Groups"](#) on page 4-42.

Overview of Oracle ASM Dynamic Volume Manager

The Oracle ASM Dynamic Volume Manager (Oracle ADVM) provides volume management services and a standard disk device driver interface to clients. File systems and other disk-based applications issue I/O requests to Oracle ADVM volume devices as they would to other storage devices on a vendor operating system.

An Oracle ADVM volume device is constructed from an Oracle ASM dynamic volume. One or more Oracle ADVM volume devices may be configured within each Oracle ASM disk group. The Oracle ADVM Driver maps I/O requests against an Oracle ADVM volume device to blocks in a corresponding Oracle ASM dynamic volume and disk set located within an Oracle ASM disk group. An Oracle ADVM volume device exports Oracle ASM volume manager features and ensures that volume mirrors remain consistent in the face of abnormal system shutdowns, Oracle ASM instance failures, or system crashes.

Oracle ADVM extends Oracle ASM by providing a disk driver interface to Oracle ASM storage allocated as Oracle ASM volume files. You can use Oracle ADVM to create virtual disks that contain file systems. These file systems contained on Oracle ASM volumes can be used to support files beyond Oracle database files, such as executable files, report files, trace files, alert logs, and other application data files. Because Oracle ADVM volumes are actually Oracle ASM files, they require the same administrative privileges as the Oracle ASM files.

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is layered on Oracle ASM through the Oracle ADVM interface. With the addition of the Oracle

ADVM, Oracle ASM becomes a complete storage solution of user data for both database and non-database file needs.

To add a volume to an Oracle ASM disk group, disk group attributes `COMPATIBLE.ASM` and `COMPATIBLE.ADVM` must be set to '11.2'. The volume name can be a maximum of 11 alphanumeric characters; dashes are not allowed. The first character must be alphabetic.

Notes:

- Dynamic volumes supersede traditional device partitioning. Each volume is individually named and may be configured for a single file system. Oracle ADVM volumes may be created on demand from Oracle ASM disk group storage and dynamically resized as required. These attributes make Oracle ADVM volumes far more flexible than physical devices and associated partitioning schemes.
 - Oracle ADVM supports all storage solutions supported for Oracle ASM with the exception of NFS and Exadata storage.
-
-

The Oracle ADVM functionality includes the following:

- The `ALTER DISKGROUP ... ADD | RESIZE | DROP | ENABLE | DISABLE | MODIFY VOLUME SQL` statements that manage Oracle ADVM volumes.

See ["Adding Volumes to Disk Groups"](#) on page 4-13.

See Also: The `ALTER DISKGROUP SQL` statement in the *Oracle Database SQL Language Reference*

- ASMCMD command-line tools for managing Oracle ASM volumes.
See ["ASMCMD Volume Management Commands"](#) on page 12-58.
- There are new `V$ASM` views display information about Oracle ASM volumes.
See ["Views Containing Oracle ACFS Information"](#) on page 6-6.
- Oracle Enterprise Manager support for managing Oracle ASM volumes.
See [Chapter 10, "Administering Oracle ACFS With Oracle Enterprise Manager"](#).
- Oracle ACFS operating system utilities for managing file systems and volume device files.
See [Chapter 13, "Oracle ACFS Command-Line Tools"](#).

Using Views to Display Information

This chapter contains information about using dynamic views to display Oracle Automatic Storage Management (Oracle ASM), Oracle Automatic Storage Management Cluster File System (Oracle ACFS), and Oracle ASM Dynamic Volume Manager (Oracle ADVM) information.

- [Views Containing Oracle ASM Disk Group Information](#)
- [Views Containing Oracle ACFS Information](#)

See Also: *Oracle Database Reference* for information about all of the V\$ASM* dynamic performance views

Views Containing Oracle ASM Disk Group Information

You can use the views in [Table 6–1](#) to obtain information about Oracle ASM disk groups.

The V\$ASM_ALIAS, V\$ASM_CLIENT, V\$ASM_DISKGROUP, V\$ASM_DISKGROUP_STAT, V\$ASM_FILE, V\$ASM_USER, V\$ASM_USERGROUP, and V\$ASM_USERGROUP_MEMBER are accessible from both the Oracle ASM and database instances.

Table 6–1 Oracle ASM Dynamic Views for Disk Group Information

View	Description
V\$ASM_ALIAS	Contains one row for every alias present in every disk group mounted by the Oracle ASM instance.
V\$ASM_ATTRIBUTE	Displays one row for each attribute defined. In addition to attributes specified by CREATE DISKGROUP and ALTER DISKGROUP statements, the view may show other attributes that are created automatically. Attributes are only displayed for disk groups where COMPATIBLE . ASM is set to 11.1 or higher.
V\$ASM_CLIENT	In an Oracle ASM instance, identifies databases using disk groups managed by the Oracle ASM instance. In a DB instance, contains information about the Oracle ASM instance if the database has any open Oracle ASM files.
V\$ASM_DISK	Contains one row for every disk discovered by the Oracle ASM instance, including disks that are not part of any disk group. This view performs disk discovery every time it is queried.
V\$ASM_DISK_IOSTAT	Displays information about disk I/O statistics for each Oracle ASM client. In a DB instance, only the rows for that instance are shown.

Table 6–1 (Cont.) Oracle ASM Dynamic Views for Disk Group Information

View	Description
V\$ASM_DISK_STAT	Contains the same columns as V\$ASM_DISK, but to reduce overhead, does not perform a discovery when it is queried. It only returns information about any disks that are part of mounted disk groups in the storage system. To see all disks, use V\$ASM_DISK instead.
V\$ASM_DISKGROUP	Describes a disk group (number, name, size related info, state, and redundancy type). This view performs disk discovery every time it is queried.
V\$ASM_DISKGROUP_STAT	Contains the same columns as V\$ASM_DISKGROUP, but to reduce overhead, does not perform a discovery when it is queried. It does not return information about any disks that are part of mounted disk groups in the storage system. To see all disks, use V\$ASM_DISKGROUP instead.
V\$ASM_FILE	Contains one row for every Oracle ASM file in every disk group mounted by the Oracle ASM instance.
V\$ASM_OPERATION	In an Oracle ASM instance, contains one row for every active Oracle ASM long running operation executing in the Oracle ASM instance. In a DB instance, contains no rows.
V\$ASM_TEMPLATE	Contains one row for every template present in every disk group mounted by the Oracle ASM instance.
V\$ASM_USER	Contains the effective operating system user names of connected database instances and names of file owners.
V\$ASM_USERGROUP	Contains the creator for each Oracle ASM File Access Control group.
V\$ASM_USERGROUP_MEMBER	Contains the members for each Oracle ASM File Access Control group.

When querying V\$ASM views, note that the value of the disk group number is not a static value. When a disk group is mounted, a disk group number is chosen. This number may change across disk group mounts. A disk group number is not recorded in any persistent structure, but the current value can be viewed in the GROUP_NUMBER column of the V\$ASM views.

An example of the use of the V\$ASM_ATTRIBUTE and V\$ASM_DISKGROUP views is shown in [Example 6–1](#). The COMPATIBLE.ASM value must be set to 11.1 or higher for the disk group to display in the V\$ASM_ATTRIBUTE view output. Attributes that are designated as read-only (Y) can only be set during disk group creation.

Example 6–1 Viewing Disk Group Attributes With V\$ASM_ATTRIBUTE

```
SQL> SELECT dg.name AS diskgroup, SUBSTR(a.name,1,18) AS name,
        SUBSTR(a.value,1,24) AS value, read_only FROM V$ASM_DISKGROUP dg,
        V$ASM_ATTRIBUTE a WHERE dg.name = 'DATA'
        AND dg.group_number = a.group_number;
```

DISKGROUP	NAME	VALUE	READ_ONLY
DATA	disk_repair_time	3.6h	N
DATA	au_size	4194304	Y
DATA	sector_size	512	Y
DATA	compatible.asm	11.2.0.0.0	N
DATA	compatible.rdbms	11.2.0.0.0	N
DATA	compatible.advm	11.2.0.0.0	N
DATA	cell.smart_scan_capable	FALSE	N

```

DATA          access_control.enabled  TRUE          N
DATA          access_control.umask    002          N
...

```

You can view the compatibility for a disk group with the V\$ASM_DISKGROUP view, as shown in [Example 6-2](#).

Example 6-2 Viewing the Compatibility of a Disk Group with V\$ASM_DISKGROUP

```

SQL> SELECT name AS diskgroup, compatibility AS asm_compat,
       database_compatibility AS db_compat FROM V$ASM_DISKGROUP;

```

```

DISKGROUP      ASM_COMPAT  DB_COMPAT
-----
DATA           11.2.0.0.0  11.2.0.0.0
FRA            10.1.0.0.0  10.1.0.0.0

```

An example of the use of the V\$ASM_DISK and V\$ASM_DISKGROUP views is shown in [Example 6-3](#). This example displays the disks associated with a disk group, plus the mount status and state of the disks.

Example 6-3 Viewing Disks in Disk Groups with V\$ASM_DISK

```

SQL> SELECT SUBSTR(d.name,1,16) AS asmdisk, d.mount_status, d.state,
       dg.name AS diskgroup FROM V$ASM_DISKGROUP dg, V$ASM_DISK d
       WHERE dg.group_number = d.group_number;

```

```

ASMDISK          MOUNT_S STATE  DISKGROUP
-----
DATA_0001        CACHED  NORMAL  DATA
DATA_0007        CACHED  NORMAL  DATA
DATA_0000        CACHED  NORMAL  DATA
DATA_0008        CACHED  NORMAL  DATA
DATA_0005        CACHED  NORMAL  DATA
DATA_0002        CACHED  NORMAL  DATA
DATA_0004        CACHED  NORMAL  DATA
DATA_0006        CACHED  NORMAL  DATA
DATA_0003        CACHED  NORMAL  DATA
FRA_0005         CACHED  NORMAL  FRA
FRA_0004         CACHED  NORMAL  FRA
FRA_0000         CACHED  NORMAL  FRA
FRA_0009         CACHED  NORMAL  FRA
FRA_0011         CACHED  NORMAL  FRA
...

```

An example of the use of the V\$ASM_CLIENT and V\$ASM_DISKGROUP views on an Oracle ASM instance is shown in [Example 6-4](#). This example displays disk groups with information about the connected database client instances.

Example 6-4 Viewing Disk Group Clients With V\$ASM_CLIENT

```

SQL> SELECT dg.name AS diskgroup, SUBSTR(c.instance_name,1,12) AS instance,
       SUBSTR(c.db_name,1,12) AS dbname, SUBSTR(c.SOFTWARE_VERSION,1,12) AS software,
       SUBSTR(c.COMPATIBLE_VERSION,1,12) AS compatible
       FROM V$ASM_DISKGROUP dg, V$ASM_CLIENT c
       WHERE dg.group_number = c.group_number;

```

```

DISKGROUP          INSTANCE  DBNAME  SOFTWARE  COMPATIBLE
-----
DATA               +ASM     +ASM    11.2.0.0.0  11.2.0.0.0

```

DATA	orcl	orcl	11.2.0.0.0	11.2.0.0.0
DATA	+ASM	asmvol	11.2.0.0.0	11.2.0.0.0
FRA	orcl	orcl	11.2.0.0.0	11.2.0.0.0
...				

Viewing Oracle ASM File Access Control Information

You can view information about Oracle ASM File Access Control in the columns of the \$ASM_USER, V\$ASM_USERGROUP, V\$ASM_USERGROUP_MEMBER, and V\$ASM_FILE views.

[Example 6-5](#) shows information about Oracle ASM File Access Control users displayed in the V\$ASM_USER view.

Example 6-5 Viewing Access Control Information in V\$ASM_USER

```
SQL> SELECT dg.name AS diskgroup, u.group_number, u.user_number, u.os_id, u.os_name
FROM V$ASM_DISKGROUP dg, V$ASM_USER u
WHERE dg.group_number = u.group_number AND dg.name = 'DATA';
```

DISKGROUP	GROUP_NUMBER	USER_NUMBER	OS_ID	OS_NAME
DATA	1	1	1001	oracle1
DATA	1	2	1002	oracle2
DATA	1	3	1003	grid

[Example 6-6](#) shows information about Oracle ASM File Access Control user groups displayed in the V\$ASM_USERGROUP view.

Example 6-6 Viewing Access Control Information With V\$ASM_USERGROUP

```
SQL> SELECT dg.name AS diskgroup, ug.group_number, ug.owner_number, u.os_name,
ug.usergroup_number, ug.name FROM V$ASM_DISKGROUP dg, V$ASM_USER u, V$ASM_USERGROUP ug
WHERE dg.group_number = ug.group_number AND dg.name = 'DATA'
AND ug.owner_number = u.user_number;
```

DISKGROUP	GROUP_NUMBER	OWNER_NUMBER	OS_NAME	USERGROUP_NUMBER	NAME
DATA	1	3	grid	1	asm_data

[Example 6-7](#) shows information about Oracle ASM File Access Control user groups and members displayed in the V\$ASM_USERGROUP_MEMBER view.

Example 6-7 Viewing Access Control Information With V\$ASM_USERGROUP_MEMBER

```
SQL> SELECT dg.name AS diskgroup, um.group_number, um.member_number, u.os_name,
um.usergroup_number, ug.name FROM V$ASM_DISKGROUP dg, V$ASM_USER u, V$ASM_USERGROUP_MEMBER um,
V$ASM_USERGROUP ug WHERE dg.group_number = um.group_number AND
dg.group_number = ug.group_number AND dg.group_number = u.group_number AND dg.name = 'DATA'
AND um.member_number = u.user_number AND um.usergroup_number = ug.usergroup_number;
```

DISKGROUP	GROUP_NUMBER	MEMBER_NUMBER	OS_NAME	USERGROUP_NUMBER	NAME
DATA	1	1	oracle1	1	asm_data
DATA	1	2	oracle2	1	asm_data

[Example 6-8](#) shows information about Oracle ASM File Access Control file permissions displayed in the V\$ASM_FILE view.

Example 6–8 Viewing Access Control Information With V\$ASM_FILE

```
SQL> SELECT dg.name AS diskgroup, a.name, f.permissions, f.user_number, u.os_name,
f.usergroup_number, ug.NAME FROM V$ASM_DISKGROUP dg, V$ASM_USER u, V$ASM_USERGROUP ug,
V$ASM_FILE f, V$ASM_ALIAS a WHERE dg.group_number = f.group_number AND
dg.group_number = u.group_number AND dg.group_number = ug.group_number AND
dg.name = 'FRA' AND f.usergroup_number = ug.usergroup_number AND f.user_number = u.user_number
AND f.file_number = a.file_number;
```

DISKGROUP	NAME	PERMISSIONS	USER_NUMBER	OS_NAME	USERGROUP_NUMBER	NAME
DATA	USERS.259.685366091	rw-r-----	3	grid	1	asm_fra
DATA	TEMP.264.685366227	rw-r-----	3	grid	1	asm_fra
...						

For more information about Oracle ASM File Access Control, see "[Managing Oracle ASM File Access Control for Disk Groups](#)" on page 4-38.

Viewing Disk Region Information

Information about Intelligent Data Placement is displayed in the columns of the \$ASM_DISK, V\$ASM_DISK_IOSTAT, V\$ASM_FILE, and V\$ASM_TEMPLATE views.

[Example 6–9](#) shows queries for Intelligent Data Placement information in the PRIMARY_REGION and MIRROR_REGION columns of the V\$ASM_FILE view.

Example 6–9 Viewing Intelligent Data Placement Information With V\$ASM_FILE

```
SQL> SELECT dg.name AS diskgroup, f.file_number, f.primary_region, f.mirror_region, f.hot_reads,
f.hot_writes, f.cold_reads, f.cold_writes
FROM V$ASM_DISKGROUP dg, V$ASM_FILE f
WHERE dg.group_number = f.group_number and dg.name = 'DATA';
```

DISKGROUP	FILE_NUMBER	PRIM	MIRR	HOT_READS	HOT_WRITES	COLD_READS	COLD_WRITES
DATA	257	COLD	COLD	0	0	119770	886575
DATA	258	COLD	COLD	0	0	1396	222282
DATA	259	COLD	COLD	0	0	2056	199
DATA	260	COLD	COLD	0	0	42377	1331016
DATA	261	COLD	COLD	0	0	4336300	1331027
...							

[Example 6–10](#) shows displays Intelligent Data Placement information in the PRIMARY_REGION and MIRROR_REGION columns of the V\$ASM_TEMPLATE view.

Example 6–10 Viewing Intelligent Data Placement Information With V\$ASM_TEMPLATE

```
SQL> SELECT dg.name AS diskgroup, t.name, t.stripe, t.redundancy, t.primary_region, t.mirror_region
FROM V$ASM_DISKGROUP dg, V$ASM_TEMPLATE t
WHERE dg.group_number = t.group_number and dg.name = 'DATA' ORDER BY t.name;
```

DISKGROUP	NAME	STRIPE	REDUND	PRIM	MIRR
DATA	ARCHIVELOG	COARSE	MIRROR	COLD	COLD
DATA	ASMPARAMETERFILE	COARSE	MIRROR	COLD	COLD
DATA	AUTOBACKUP	COARSE	MIRROR	COLD	COLD
DATA	BACKUPSET	COARSE	MIRROR	COLD	COLD
DATA	CHANGETRACKING	COARSE	MIRROR	COLD	COLD
DATA	CONTROLFILE	FINE	HIGH	COLD	COLD
DATA	DATAFILE	COARSE	MIRROR	COLD	COLD
DATA	DATAGUARDCONFIG	COARSE	MIRROR	COLD	COLD
DATA	DUMPSET	COARSE	MIRROR	COLD	COLD
DATA	FLASHBACK	COARSE	MIRROR	COLD	COLD

DATA	OCRFILE	COARSE MIRROR COLD COLD
DATA	ONLINELOG	COARSE MIRROR COLD COLD
DATA	PARAMETERFILE	COARSE MIRROR COLD COLD
DATA	TEMPFILE	COARSE MIRROR COLD COLD
DATA	XTRANSPORT	COARSE MIRROR COLD COLD

15 rows selected.

For information about setting Intelligent Data Placement, see ["Intelligent Data Placement"](#) on page 4-15.

Views Containing Oracle ACFS Information

You can use the views in [Table 6–2](#) to obtain information about Oracle Automatic Storage Management Cluster File System (Oracle ACFS). These views are accessible from the Oracle ASM instance.

Table 6–2 Dynamic View for Oracle ACFS Information

View	Description
V\$ASM_ACFSSNAPSHOTS	Contains snapshot information for every mounted Oracle ACFS file system.
V\$ASM_ACFSVOLUMES	Contains information about mounted Oracle ACFS volumes, correlated with V\$ASM_FILESYSTEM.
V\$ASM_FILESYSTEM	Contains columns that display information for every mounted Oracle ACFS file system.
V\$ASM_VOLUME	Contains information about each Oracle ADVM volume that is a member of an Oracle ASM instance.
V\$ASM_VOLUME_STAT	Contains information about statistics for each Oracle ADVM volume.

[Example 6–11](#) shows information displayed from the V\$ASM_ACFSSNAPSHOTS view. The FS_NAME column contains the mount point. The VOL_DEVICE contains the name of the Oracle ADVM device.

Example 6–11 Viewing Snapshot Information in the V\$ASM_ACFSSNAPSHOTS View

```
SELECT SUBSTR(fs_name,1,34) FILESYSTEM, SUBSTR(vol_device,1,24) DEVICE,
       SUBSTR(snap_name,1,28) SNAPSHOT, create_time TIME FROM V$ASM_ACFSSNAPSHOTS;
```

FILESYSTEM	DEVICE	SNAPSHOT	TIME
/u01/app/acfsmounts/acfs1	/dev/asm/volume1-228	mynsnapshot_0900609a	09-JUL-09
/u01/app/acfsmounts/acfs1	/dev/asm/volume1-228	mynsnapshot_0900610a	10-JUL-09
/u01/app/acfsmounts/acfs1	/dev/asm/volume1-228	mynsnapshot_0900609c	09-JUL-09
/u01/app/acfsmounts/acfs1	/dev/asm/volume1-228	mynsnapshot_0900610b	10-JUL-09

See Also: *Oracle Database Reference* for information about the V\$ASM_ACFSSNAPSHOTS view

[Example 6–12](#) shows information displayed from the V\$ASM_VACFSVOLUMES view. The PRIMARY_VOL column contains TRUE if the volume is the primary volume for the file system.

Example 6–12 Viewing Volume Information With V\$ASM_ACFSVOLUMES

```
SELECT fs_name, vol_device, primary_vol, total_mb, free_mb FROM V$ASM_ACFSVOLUMES;
```

FS_NAME	VOL_DEVICE	PRIMARY_VOL	TOTAL_MB	FREE_MB
/u01/app/acfsmounts/acfs1	/dev/asm/volume1-228	TRUE	1024000	578626.522
/u01/app/acfsmounts/acfs2	/dev/asm/volume2-375	TRUE	1024000	685761.463
...				

See Also: *Oracle Database Reference* for information about the V\$ASM_ACFSVOLUMES view

[Example 6-13](#) shows information displayed from the V\$ASM_FILESYSTEM view.

The STATE column contains the status of the file system, either AVAILABLE or OFFLINE. An offline file system can only be dismounted; other attempts at access result in errors. Offline means that either the Oracle ASM instance is down, the disk group has been forced dismounted, or less commonly, a metadata I/O failure occurred or serious metadata corruption was detected. In the case of a metadata I/O failure, the file system is also marked as corrupt.

The CORRUPT column indicates whether the file system needs the fsck or acfschkdisk command run on it. See "[fsck](#)" on page 13-7 and "[acfschkdisk](#)" on page 13-13.

Example 6-13 Viewing Volume Information with V\$ASM_FILESYSTEM

```
SQL> SELECT fs_name, available_time, block_size, state, corrupt FROM V$ASM_FILESYSTEM;
```

FS_NAME	AVAILABLE	BLOCK_SIZE	STATE	CORRUPT
/u01/app/acfsmounts/acfs1	19-JUL-09	4	AVAILABLE	FALSE
/u01/app/acfsmounts/acfs2	19-JUL-09	4	AVAILABLE	FALSE

See Also: *Oracle Database Reference* for information about the V\$ASM_FILESYSTEM view

[Example 6-14](#) shows information displayed from the V\$ASM_VOLUME view for volumes contained in the DATA disk group.

Example 6-14 Viewing Volume Information With V\$ASM_VOLUME

```
SELECT dg.name AS diskgroup, v.volume_name, v.volume_device, v.mountpath
FROM V$ASM_DISKGROUP dg, V$ASM_VOLUME v
WHERE dg.group_number = v.group_number and dg.name = 'DATA';
```

DISKGROUP	VOLUME_NAME	VOLUME_DEVICE	MOUNTPATH
DATA	VOLUME1	/dev/asm/volume1-228	/u01/app/acfsmounts/acfs1
DATA	VOLUME2	/dev/asm/volume2-375	/u01/app/acfsmounts/acfs2

See Also: *Oracle Database Reference* for information about the V\$ASM_VOLUME view

[Example 6-14](#) shows information displayed from the V\$ASM_VOLUME_STAT view for volumes contained in the DATA disk group. The BYTES_READ column contains the total number of bytes read for the volume. The BYTES_WRITTEN column contains the total number of bytes written for the volume.

Example 6–15 Viewing Volume Information With V\$ASM_VOLUME_STAT

```
SELECT dg.name AS diskgroup, v.volume_name, v.bytes_read, v.bytes_written
FROM V$ASM_DISKGROUP dg, V$ASM_VOLUME_STAT v
WHERE dg.group_number = v.group_number and dg.name = 'DATA';
```

DISKGROUP	VOLUME_NAME	BYTES_READ	BYTES_WRITTEN
DATA	VOLUME1	12370105856	43510272
DATA	VOLUME2	2685728	32201504

See Also: *Oracle Database Reference* for information about the V\$ASM_VOLUME_STAT view

Administering Oracle ASM Files, Directories, and Templates

This chapter describes how to administer files when you use the Oracle Automatic Storage Management (Oracle ASM) storage option.

This chapter contains the following topics:

- [What Types of Files Does Oracle ASM Support?](#)
- [About Oracle ASM Filenames](#)
- [Creating and Referencing Oracle ASM Files in the Database](#)
- [Managing Alias Names for Oracle ASM Filenames](#)
- [Managing Disk Group Directories](#)
- [Accessing Oracle ASM Files with the XML DB Virtual Folder](#)
- [Using DBMS_FILE Transfer Utility for Oracle ASM](#)
- [Managing Disk Group Templates](#)

What Types of Files Does Oracle ASM Support?

Oracle ASM supports most file types required by the database. [Table 7-1](#) lists file types that are supported directly by Oracle ASM and lists the system default template that provides the attributes for file creation. Some of the file types shown in the table are related to specific products or features, and are not discussed in this book.

Table 7-1 File Types Supported by Automatic Storage Management

File Type	Default Templates
Control files	CONTROLFILE
Data files	DATAFILE
Redo log files	ONLINELOG
Archive log files	ARCHIVELOG
Temporary files	TEMPFILE
Data file backup pieces	BACKUPSET
Data file incremental backup pieces	BACKUPSET
Archive log backup piece	BACKUPSET
Data file copy	DATAFILE

Table 7-1 (Cont.) File Types Supported by Automatic Storage Management

File Type	Default Templates
Persistent initialization parameter file (SPFILE)	PARAMETERFILE
Flashback logs	FLASHBACK
Change tracking file	CHANGETRACKING
Data Pump dumpset	DUMPSET
Automatically generated control file backup	AUTOBACKUP
Cross-platform transportable data files	XTRANSPORT
Flash file	FLASHFILE
Oracle ASM Persistent initialization parameter file (SPFILE)	ASMPARAMETERFILE
Oracle ASM Persistent initialization parameter file (SPFILE) backup	ASMPARAMETERFILEBACKUP
Oracle Cluster Registry file	OCRFILE
Oracle ASM Dynamic Volume Manager volumes	n/a

Oracle ASM cannot directly support some administrative file types on disk groups. These include trace files, audit files, alert logs, export files, tar files, and core files.

Note: Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) extend Oracle ASM support to include database and application executables, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data. See [Chapter 5, "Introduction to Oracle ACFS"](#).

See Also: ["Managing Disk Group Templates"](#) on page 7-15 for a description of the system default templates

About Oracle ASM Filenames

Every file created in Oracle ASM is provided a filename generated by Oracle-Managed Files, known as a fully qualified filename (or system-generated filename). The fully qualified filename represents a complete path name in the Oracle ASM file system. An example of a fully qualified filename is:

```
+data/orcl/controlfile/Current.256.541956473
```

You can use the fully qualified filename to reference (read or retrieve) an Oracle ASM file. You can also use other abbreviated filename formats, such as an alias Oracle ASM filename described in ["Alias Oracle ASM Filename Forms"](#) on page 7-5, to reference an Oracle ASM file.

Oracle ASM generates a fully qualified filename upon any request to create a file. A creation request does not or cannot specify a fully qualified filename. Instead, it uses a simpler syntax to specify a file, such as an alias or just a disk group name. Oracle ASM then creates the file, placing it in the correct Oracle ASM path according to file type, and then assigns an appropriate fully qualified filename. If you specify an alias in the

creation request, Oracle ASM also creates the alias so that it references the fully qualified filename.

Oracle ASM file creation requests are either single file creation requests or multiple file creation request.

Notes:

- You can find the generated name in database views displaying Oracle file names, such as V\$DATAFILE and V\$LOGFILE. You can use this name, or an abbreviated form of it, if you later need to reference an Oracle ASM file in a SQL statement. Like other Oracle Database filenames, Oracle ASM filenames are kept in the control file and the RMAN catalog.
 - Fully qualified and numeric filenames can be used in single-file create if you specify the REUSE keyword, as described in "Using Oracle ASM Filenames in SQL Statements" on page 7-9.
 - Forward slash (/) and backslash (\) are interchangeable in filenames. Filenames are case insensitive, but case retentive.
-
-

This sections contains the following topics:

- [Single File Creation Form](#)
- [Multiple File Creation Form](#)
- [Fully qualified File Name Form](#)
- [Alias Oracle ASM Filename Forms](#)
- [Alias Oracle ASM Filename with Template Form](#)
- [Incomplete Oracle ASM Filename Form](#)
- [Incomplete Oracle ASM Filename with Template Form](#)

[Table 7-2](#) specifies the valid usage for each filename form, and if the form is used for file creation, whether the created file is an Oracle Managed Files (OMF).

Table 7-2 Valid Usage for Oracle ASM Filename Forms

Filename Form	Reference to File?	Single-file Creation?	Multiple File Creation?	Created as OMF?
Fully qualified file name	Yes	No	No	No
Alias file name	Yes	Yes	No	No
Alias with template file name	No	Yes	No	No
Incomplete file name	No	Yes	Yes	Yes
Incomplete file name with template	No	Yes	Yes	Yes

Single File Creation Form

A single file creation form is used to create a single file, such as a data file or a control file. The form of the Oracle ASM filename in this type is either an alias, such as +data/orcl/ctl.f, or an incomplete file name, such as a disk group name preceded by a plus sign (+data). You use the alias or disk group name where a

filename is called for in an SQL statement, such as `CREATE TABLESPACE`, `CREATE CONTROLFILE`.

Multiple File Creation Form

A multiple file creation form is used multiple times to create Oracle ASM files, such as when assigning a value to the initialization parameter `DB_CREATE_FILE_DEST`. You can also issue a `CREATE TABLESPACE` statement (without a filename specification) multiple times. Each time, Oracle ASM creates a different unique data file name.

When using the multiple file creation form, you can use an incomplete filename form, such as a disk group name preceded by a plus sign (`+data`). In this case, you set `DB_CREATE_FILE_DEST` to an incomplete filename (for example, `+data`), and whenever a command is executed that must create a database file in `DB_CREATE_FILE_DEST`, the file is created in the designated disk group and assigned a unique fully qualified name. You can use an incomplete filename in other `*_DEST` initialization parameters.

Fully qualified File Name Form

A fully qualified form can be used for referencing existing Oracle ASM files in Oracle ASM operations, except for disk group creation. This form is the Oracle ASM file name that Oracle ASM automatically generates when an Oracle ASM file is created.

A fully qualified file name has the following form:

```
+diskgroup/dbname/filetype/filetypetag.file.incarnation
```

Where:

- *+diskgroup* is the disk group name preceded by a plus sign.
You can think of the plus sign (+) as the root directory of the Oracle ASM file system, similar to the slash (/) on UNIX or Linux computers.
- *dbname* is the `DB_UNIQUE_NAME` of the database to which the file belongs.
- *filetype* is the Oracle file type and can be one of the file types shown in [Table 7-3](#).
- *filetypetag* is type specific information about the file and can be one of the tags shown in [Table 7-3](#).
- *file.incarnation* is the file/incarnation pair, used to ensure uniqueness.

An example of a fully qualified Oracle ASM filename is:

```
+data/orcl/controlfile/Current.256.541956473
```

Table 7-3 Oracle File Types and Oracle ASM File Type Tags

Oracle ASM File Type	Description	Oracle ASM File Type Tag	Comments
CONTROLFILE	Control files and backup control files	Current	Archive
		Backup	Backup
DATAFILE	Data files and data file copies	<i>tblspname</i>	Name of the tablespace into which the file is added
ONLINELOG	Online logs	<i>group_loggroup#</i>	
ARCHIVELOG	Archive logs	<i>thread_thread#_seq_sequence#</i>	

Table 7–3 (Cont.) Oracle File Types and Oracle ASM File Type Tags

Oracle ASM File Type	Description	Oracle ASM File Type Tag	Comments
TEMPFILE	Tempfiles	<i>tname</i>	Name of the tablespace into which the file is added
BACKUPSET	Data file and archive log backup pieces; data file incremental backup pieces	<i>spfile?_timestamp</i>	<i>spfile?</i> can take one of two values: <i>s</i> indicates that the backup set includes the <i>spfile</i> ; <i>n</i> indicates that the backup set does not include the <i>spfile</i> . <i>timestamp</i> is the time that the backup is started.
PARAMETERFILE	Persistent parameter files	<i>spfile</i>	Name of the SPFILE
DATAGUARDCONFIG	Data Guard configuration file	<i>db-unique-name</i>	Data Guard uses the value of the DB_UNIQUE_NAME initialization parameter.
FLASHBACK	Flashback logs	<i>log_flashbacklog#</i>	
CHANGETRACKING	Block change tracking data	<i>ctf</i>	Used during incremental backups
DUMPSET	Data Pump dumpset	<i>user_obj#_file#</i>	Dump set files encode the user name, the job number that created the dump set, and the file number as part of the tag.
XTRANSPORT	Data file convert	<i>tblspname</i>	Name of the tablespace into the Oracle ASM metadata is added
AUTOBACKUP	Automatic backup files	<i>spfile?_timestamp</i>	<i>spfile?</i> can take one of two values: <i>s</i> indicates that the backup set includes the <i>spfile</i> ; <i>n</i> indicates that the backup set does not include the <i>spfile</i> . <i>timestamp</i> is the time that the backup is started.
ASMPARAMETERFILE	Oracle ASM persistent parameter files	<i>spfile</i>	Name of the Oracle ASM SPFILE
OCRFILE	Oracle Clusterware Registry files	<i>ocrfile</i>	Name of the OCR files

Alias Oracle ASM Filename Forms

Alias Oracle ASM filenames can be used both for referencing existing Oracle ASM files and for creating new Oracle ASM files. Alias names start with the disk group name preceded by a plus sign, after which you specify a name string. Alias filenames are implemented using a hierarchical directory structure, with the slash (/) or backslash (\) character separating name components. You can create an alias in any Oracle ASM directory generated by Oracle-Managed Files (system-generated) or created by a user. An alias must include the disk group name; you cannot create an alias at the root level (+).

When you create an Oracle ASM file with an alias filename, the file is created with a fully qualified name, and the alias filename is additionally created. You can then access the file with either name.

Alias Oracle ASM filenames are distinguished from fully qualified filenames or numeric filenames because they do not end in a dotted pair of numbers. It is an error

to attempt to create an alias that ends in a dotted pair of numbers, such as in the format `USERS.259.685366091`. Examples of Oracle ASM alias filenames are:

```
+data/orcl/control_file1  
+fra/recover/second.dbf
```

Oracle Database references database files by their alias filenames, but only if you create the database files with aliases. If you create database files without aliases and then add aliases later, the database references the files by their fully qualified filenames. The following are examples of how the database uses alias filenames:

- Alias filenames appear in `V$` views. For example, if you create a tablespace and use an alias filename for the data file, the `V$DATAFILE` view shows the alias filename.
- When a control file points to data files and online redo log files, it can use alias filenames.
- The `CONTROL_FILES` initialization parameter can use the alias filenames of the control files. The Database Configuration Assistant (DBCA) creates control files with alias filenames.

Note: Files created using an alias filename are not considered Oracle Managed Files and might require manual deletion in the future if they are no longer needed.

For more information about alias filenames, see "[Managing Alias Names for Oracle ASM Filenames](#)" on page 7-9.

Creating a Tablespace in Oracle ASM: Using a Data File with an Alias Name

[Example 7-1](#) creates an undo tablespace with a data file that has an alias name, and with attributes that are set by the user-defined template `my_undo_template`. This example assumes that a directory has been created in disk group `data` to contain the alias name and that the user-defined template exists.

Example 7-1 Creating a Tablespace Using an Alias Name

```
CREATE UNDO TABLESPACE myundo  
    DATAFILE '+data(my_undo_template)/orcl/my_undo_ts' SIZE 200M;
```

Because an alias is used to create the data file, the file is not an Oracle Managed Files (OMF) file and the file is not be automatically deleted when the tablespace is dropped. To drop the file manually after the tablespace has been dropped, use the following SQL statement:

```
ALTER DISKGROUP data DROP FILE '+data/orcl/my_undo_ts';
```

Alias Oracle ASM Filename with Template Form

An alias Oracle ASM filename with template is used only for Oracle ASM file creation operations. It has the following format:

```
+diskgroup(template)/alias
```

Alias filenames with template behave identically to alias filenames except that a file created with an alias filename with template receives the attributes specified by the named template. The template must belong to the disk group that the file is being created in. See [Example 7-1](#) on page 7-6.

The creation and maintenance of Oracle ASM templates is discussed in ["Managing Disk Group Templates"](#) on page 7-15.

An example of an alias Oracle ASM filename with template is:

```
+data(mytemplate)/config1
```

Explicitly specifying a template name, as in the previous example, overrides the system default template for the type of file being created.

Note: Files created using an alias filename with template are not considered Oracle Managed Files and might require manual deletion in the future if they are no longer needed.

Incomplete Oracle ASM Filename Form

Incomplete Oracle ASM filenames are used only for file creation operations and are used for both single and multiple file creation. They consist only of the disk group name. Oracle ASM uses a system default template to determine the Oracle ASM file mirroring and striping attributes. The system template that is used is determined by the file type that is being created. For example, if you are creating a data file for a tablespace, the `DATAFILE` template is used.

An example of using an incomplete Oracle ASM filename is setting the `DB_CREATE_FILE_DEST` initialization parameter to:

```
+data
```

With this setting, every time you create a tablespace, a data file is created in the disk group `data`, and each data file is assigned a different fully qualified name. See ["Creating Oracle ASM Files Using a Default File Location for Disk Group Specification"](#) on page 7-8 for more information.

Incomplete Oracle ASM Filename with Template Form

Incomplete Oracle ASM filenames with templates are used only for file creation operations and are used for both single and multiple file creation. They consist of the disk group name followed by the template name in parentheses. When you explicitly specify a template in a file name, Oracle ASM uses the specified template instead of the default template for that file type to determine attributes for the file.

An example of using an incomplete Oracle ASM filename with template is setting the `DB_CREATE_FILE_DEST` initialization parameter to:

```
+data(mytemplate)
```

Creating and Referencing Oracle ASM Files in the Database

An Oracle ASM file is an Oracle Managed File unless you created the file using an alias. Any Oracle Managed File is automatically deleted when it is no longer needed.

This section contains the following topics:

- [Creating Oracle ASM Files Using a Default File Location for Disk Group Specification](#)
- [Using Oracle ASM Filenames in SQL Statements](#)

Creating Oracle ASM Files Using a Default File Location for Disk Group Specification

Using the Oracle Managed Files feature for operating system files, you can specify a directory as the default location for the creation of data files, temporary files, redo log files, and control files. Using the Oracle Managed Files feature for Oracle ASM, you can specify a disk group, in the form of an incomplete Oracle ASM filename, as the default location for creation of these files, and additional types of files, including archived log files. As for operating system files, the name of the default disk group is stored in an initialization parameter and is used whenever a file specification (for example, `DATAFILE` clause) is not explicitly specified during file creation.

The following initialization parameters accept the multiple file creation context form of Oracle ASM filenames as a destination:

- `DB_CREATE_FILE_DEST`
Specifies the default disk group location in which to create data and temp files.
- `DB_CREATE_ONLINE_LOG_DEST_n`
Specifies the default disk group location in which to create redo log and control files.
- `DB_RECOVERY_FILE_DEST`
Specifies a default disk group for a fast recovery area that contains a copy of the control file and redo log files if this parameter is specified and `DB_CREATE_ONLINE_LOG_DEST_n` and `CONTROL_FILES` are not specified.
- `CONTROL_FILES`
Specifies a disk group in which to create control files.

The following initialization parameters accept the multiple file creation context form of the Oracle ASM filenames and Oracle ASM directory names as a destination:

- `LOG_ARCHIVE_DEST_n`
Specifies a default disk group or Oracle ASM directory as destination for archiving redo log files
- `LOG_ARCHIVE_DEST`
Optional parameter to use to specify a default disk group or Oracle ASM directory as destination for archiving redo log files. Use when specifying only one destination.
- `STANDBY_ARCHIVE_DEST`
Relevant only for a standby database. It specifies a default disk group or Oracle ASM directory that is the location of archive logs arriving from a primary database. Not discussed in this book. This parameter has been deprecated. See *Oracle Data Guard Concepts and Administration*.

The following `CREATE TABLESPACE` SQL statement illustrates how an Oracle ASM file, in this case a data file, might be created in a default disk group. This example assumes that the `DB_CREATE_FILE_DEST` initialization parameter has been set to `+data`.

```
SQL> CREATE TABLESPACE mytblspace;
```

Oracle ASM automatically creates and manages the data file for `mytblspace` on Oracle ASM disks in the disk group `data`. File extents are stored using the attributes defined by the default template for a data file.

See Also:

- *Oracle Database Reference* for information about initialization parameters
- *Oracle Database Administrator's Guide* for information about initialization parameter files

Using Oracle ASM Filenames in SQL Statements

You can specify Oracle ASM filenames in the file specification clause of your SQL statements. If you are creating a file for the first time, use the creation form of an Oracle ASM filename. If the Oracle ASM file exists, you must use the reference form of the filename, and if you are trying to re-create the file, you must add the `REUSE` keyword so that the space is reused for the new file. This usage might occur when, for example, trying to re-create a control file, as shown in ["Creating Control Files in Oracle ASM"](#) on page A-3.

If a reference form is used with the `REUSE` keyword and the file does not exist, an error results.

[Example 7-2](#) an example of specifying an Oracle ASM filename in a SQL statement. In this case, it is used in the file creation form.

Example 7-2 Using an Oracle ASM File Name in a SQL Statement

```
CREATE TABLESPACE mytblspace DATAFILE '+data' SIZE 200M AUTOEXTEND ON;
```

The tablespace `mytblspace` is created and is comprised of one data file of size 200 MB contained in the disk group `data`. The data file is set to auto-extensible with an unlimited maximum size. An `AUTOEXTEND` clause can be used to override this default.

For more examples of the use of Oracle ASM filename in SQL statements, see [Appendix A, "Creating Database Files in Oracle ASM Using SQL*Plus"](#).

Managing Alias Names for Oracle ASM Filenames

Alias names, or aliases, are intended to provide a more user-friendly means of referring to Oracle ASM files, rather than using the system-generated filenames.

You can create an alias for a file when you create it in the database, or you can add an alias to an existing file using the `ADD ALIAS` clause of the `ALTER DISKGROUP` statement. You can create an alias in any system-generated or user-created Oracle ASM directory. An alias must include the disk group name; you cannot create an alias at the root level (+).

For more information about creating aliases with `ASMCMD`, refer to ["mkalias"](#) on page 12-25.

This section contains the following sections:

- [Adding an Alias Name for an Oracle ASM Filename](#)
- [Renaming an Alias Name for an Oracle ASM Filename](#)
- [Dropping an Alias Name for an Oracle ASM Filename](#)
- [Dropping Files and Associated Aliases from a Disk Group](#)

Adding an Alias Name for an Oracle ASM Filename

Use the `ADD ALIAS` clause of the `ALTER DISKGROUP` statement to create an alias name for an Oracle ASM filename. The alias name must consist of the full directory path and the alias itself.

[Example 7-3](#) adds a new alias name for a system-generated file name.

Example 7-3 Adding an Alias Name for an Oracle ASM Filename

```
ALTER DISKGROUP data ADD ALIAS '+data/orcl/second.dbf'  
FOR '+data/orcl/datafile/mytable.342.123456789';
```

Renaming an Alias Name for an Oracle ASM Filename

Use the `RENAME ALIAS` clause of the `ALTER DISKGROUP` statement to rename an alias for an Oracle ASM filename. The old and the new alias names must consist of the full directory paths of the alias names.

[Example 7-4](#) renames an alias.

Example 7-4 Example: Renaming an Alias Name for an Oracle ASM Filename

```
ALTER DISKGROUP data RENAME ALIAS '+data/orcl/datafile.dbf'  
TO '+data/payroll/compensation.dbf';
```

Renaming an alias does not update the database reference to the alias.

Dropping an Alias Name for an Oracle ASM Filename

Use the `DROP ALIAS` clause of the `ALTER DISKGROUP` statement to drop an alias for an Oracle ASM filename. The alias name must consist of the full directory path and the alias itself. The underlying file to which the alias refers is unchanged.

[Example 7-5](#) drops an alias.

Example 7-5 Dropping an Alias Name for an Oracle ASM Filename

```
ALTER DISKGROUP data DROP ALIAS '+data/payroll/compensation.dbf';
```

[Example 7-6](#) fails because it attempts to drop a system-generated filename and this syntax is not allowed.

Example 7-6 Dropping an Alias Name for an Oracle ASM Filename

```
ALTER DISKGROUP data  
DROP ALIAS '+data/orcl/datafile/mytable.342.123456789';
```

Dropping Files and Associated Aliases from a Disk Group

You can delete Oracle ASM files and their associated aliases from a disk group using the `DROP FILE` clause of the `ALTER DISKGROUP` statement. You must use a fully qualified filename, or an alias name when specifying the file that you want to delete. The underlying file on the file system is not dropped when the alias is dropped.

Some reasons why you might need to delete files include:

- A file created using aliases is not an Oracle Managed File. Consequently, it is not automatically deleted.

- A point-in-time-recovery of a database might restore the database to a time before a tablespace was created. The restore does not delete the tablespace, but there is no reference to the tablespace or its data file in the restored database. You could manually delete the data file.

In [Example 7-7](#), the alias name for the file is used to delete both the file and the alias from a disk group.

Example 7-7 Dropping Files and Associated Aliases from a Disk Group

```
ALTER DISKGROUP data DROP FILE '+data/payroll/compensation.dbf';
```

In [Example 7-8](#), the Oracle-Managed Files (system-generated) file name is used to drop the file and any associated alias.

Example 7-8 Dropping Files and Associated Aliases from a Disk Group

```
ALTER DISKGROUP data
  DROP FILE '+data/orcl/datafile/mytable.342.123456789';
```

Managing Disk Group Directories

Oracle ASM disk groups contain a system-generated hierarchical directory structure for storing Oracle ASM files. The system-generated filename that Oracle ASM assigns to each file represents a path in this directory hierarchy. The following is an example of a system-generated filename:

```
+data/orcl/CONTROLFILE/Current.256.541956473
```

The plus sign represents the root of the Oracle ASM file system. The `data` directory is the parent directory for all files in the `data` disk group. The `orcl` directory is the parent directory for all files in the `orcl` database, and the `CONTROLFILE` directory contains all control files for the `orcl` database.

You can create your own directories within this hierarchy to store aliases that you create. Thus, in addition to having user-friendly alias names for Oracle ASM files, you can have user-friendly paths to those names.

This section describes how to use the `ALTER DISKGROUP` statement to create a directory structure for aliases. It also describes how you can rename a directory or drop a directory. This section contains the following topics:

- [Creating a New Directory](#)
- [Renaming a Directory](#)
- [Dropping a Directory](#)

Creating a New Directory

Use the `ADD DIRECTORY` clause of the `ALTER DISKGROUP` statement to create a hierarchical directory structure for alias names for Oracle ASM files. Use the slash (/) or backslash (\) character to separate components of the directory path. The directory path must start with the disk group name, preceded by a plus sign (+), followed by any subdirectory names of your choice.

The parent directory must exist before attempting to create a subdirectory or alias in that directory. A directory must be created below the disk group level.

[Example 7-9](#) creates a hierarchical directory for disk group `data`, which can contain, for example, the alias name `+data/orcl/control_file1`.

Example 7–9 Creating a New Directory

```
ALTER DISKGROUP data ADD DIRECTORY '+data/orcl';
```

Assuming no subdirectory exists under the directory `+data/orcl`, the SQL statement fails in [Example 7–10](#).

Example 7–10 Creating a New Subdirectory

```
ALTER DISKGROUP data
  ADD DIRECTORY '+data/orcl/first_dir/second_dir';
```

Renaming a Directory

The `RENAME DIRECTORY` clause of the `ALTER DISKGROUP` statement enables you to rename a directory. System-created directories (those containing system-generated names) cannot be renamed. The root level (disk group name) cannot be renamed.

[Example 7–11](#) renames a directory.

Example 7–11 Renaming a Directory

```
ALTER DISKGROUP data RENAME DIRECTORY '+data/mydir'
  TO '+data/yourdir';
```

Dropping a Directory

You can delete a directory using the `DROP DIRECTORY` clause of the `ALTER DISKGROUP` statement. You cannot drop a system-created directory. You cannot drop a directory containing alias names unless you also specify the `FORCE` clause.

[Example 7–12](#) deletes a directory along with its contents.

Example 7–12 Dropping a Directory

```
ALTER DISKGROUP data DROP DIRECTORY '+data/yourdir' FORCE;
```

Accessing Oracle ASM Files with the XML DB Virtual Folder

Oracle ASM files and directories can be accessed through a virtual folder in the XML DB repository. The repository path to the virtual folder is `/sys/asm`. The folder is virtual because its contents do not actually reside in the repository; they exist as normal Oracle ASM files and directories. `/sys/asm` provides a means to access and manipulate the Oracle ASM files and directories with programmatic APIs such as the `DBMS_XDB` package and with XML DB protocols such as FTP and HTTP/WebDAV. You can also use the `ASMCMD cp` command to copy files. See "[cp](#)" on page 12-19.

A typical use for this capability might be to view `/sys/asm` as a Web Folder in a graphical user interface (with the WebDAV protocol), and then copy a Data Pump dumpset from an Oracle ASM disk group to an operating system file system by dragging and dropping.

You must log in as a user other than `SYS` and you must have been granted the `DBA` role to access `/sys/asm` with XML DB protocols.

Note: For security reasons, FTP is disabled, by default. This is because the IETF FTP protocol specification requires that passwords be transmitted in clear text. Disabling is done by configuring the FTP server port as zero (0). To enable FTP, set the `ftp-port` parameter to the FTP port to use, such as 2100.

See Also:


- *Oracle XML DB Developer's Guide* for information about Oracle XML DB, including additional ways to configure port numbers for the XML DB protocol servers
- *Oracle Database PL/SQL Packages and Types Reference* for information about the `DBMS_XDB` package

Inside `/sys/asm`

The Oracle ASM virtual folder is created by default during XML DB installation. If the database is not configured to use Oracle ASM, the folder is empty and no operations are permitted on it.

The Oracle ASM virtual folder contains folders and subfolders that follow the hierarchy defined by the structure of an Oracle ASM fully qualified file name. [Figure 7-1](#) illustrates an example of this hierarchy, which for simplicity, excludes directories created for aliases.

The folder `/sys/asm` contains one subfolder for every mounted disk group, and each disk group folder contains one subfolder for each database that uses the disk group. In addition, a disk group folder might contain files and folders corresponding to aliases created by the administrator. Continuing the hierarchy, the database folders contain file type folders, which contain the Oracle ASM files.

Figure 7–1 Hierarchical Structure of Oracle ASM Folders in an XML DB Installation

The following are usage restrictions on `/sys/asm`:

- You cannot create hard links to existing Oracle ASM files or directories with APIs such as `DBMS_XDB.LINK`.
- You cannot rename (move) an Oracle ASM file to another disk group or to a directory outside Oracle ASM.

You can use the directory `/sys/asm` for storing the names of disk groups. You cannot store other files in this directory. Within the disk group directories under `/sys/asm`, such as `/sys/asm/DATA`, you can only store database files in these sub-directories. Oracle ASM rejects attempts to store non-database files in these directories.

Using DBMS_FILE Transfer Utility for Oracle ASM

The `DBMS_FILE_TRANSFER` package provides procedures to copy Oracle ASM files within a database or to transfer binary files between databases that use Oracle ASM. The `DBMS_FILE_TRANSFER` package has the following procedures:

- `COPY_FILE`—Reads a file from a source directory and creates a copy of the file in a destination directory. The source and destination directories can both be in a local file system or in an Oracle ASM disk group. You can also use this procedure to copy between a local file system and an Oracle ASM disk group; the copy operation is valid in either direction.
- `GET_FILE`—Contacts a remote database to read a remote file and then creates a copy of the file in the local file system or Oracle ASM disk group.
- `PUT_FILE`—Reads a local file or Oracle ASM disk group and contacts a remote database to create a copy of the file in the remote file system.

See Also: *Oracle Database PL/SQL Packages and Types Reference* for more information about the `DBMS_FILE_TRANSFER` package

Managing Disk Group Templates

This section describes how to manage disk group templates under the following topics:

- [Template Attributes](#)
- [Adding Templates to a Disk Group](#)
- [Modifying a Disk Group Template](#)
- [Dropping Templates from a Disk Group](#)
- [Creating Tablespaces in Oracle ASM: Specifying Attributes with Templates](#)

Templates are used to set attributes of files created in an Oracle ASM disk group. When a file is created, redundancy and striping attributes are set for that file based on an explicitly named template or the system template that is the default template for the file type.

When a disk group is created, Oracle ASM creates a set of default templates for that disk group. The set consists of one template for each file type (data file, control file, redo log file, and so on) that is supported by Oracle ASM. For example, a template named `ONLINELOG` provides the default file redundancy and striping attributes for all redo log files written to Oracle ASM disks. Default template settings depend on the disk group type. The default template for data files for a normal redundancy disk group sets two-way mirroring, while the corresponding default template in a high redundancy disk group sets three-way mirroring. You can modify these default templates.

For example, default redundancy for the online redo log files (`ONLINELOG` template) for a normal redundancy disk group is `MIRROR`. In [Example 4-1](#) on page 4-5, this means that when one copy of a redo log file extent is written to a disk in failure group `controller1`, a mirrored copy of the file extent is written to a disk in failure group `controller2`. To support the default mirroring of a normal redundancy disk group, at least two failure groups must be defined.

[Table 7-6](#) on page 7-16 lists the default templates and the attributes that are associated to matching files. As the table shows, the initial redundancy value of each default template depends on the type of disk group that the template belongs to.

Using clauses of the `ALTER DISKGROUP` statement, you can add new templates to a disk group, modify existing ones, or drop templates. The reason to add templates is to create the right combination of attributes to meet unique requirements. You can then reference a template name when creating a file, thereby assigning desired attributes based on an individual file rather than on the file type.

The `V$ASM_TEMPLATE` view lists all of the templates known to the Oracle ASM instance. For an example of the information displayed in the `V$ASM_TEMPLATE` view, see [Example 6-10](#) on page 6-5.

Template Attributes

[Table 7-4](#) shows the permitted striping attribute values and allocation unit size chunks. These values correspond to the `STRIPE` column of `V$ASM_TEMPLATE`.

Table 7-4 Permitted Values for Oracle ASM Template Striping Attribute

Striping Attribute Value	Description
FINE	Striping in 128 KB chunks.
COARSE	Striping in 1 MB chunks.

Table 7-5 shows the permitted redundancy values for Oracle ASM templates. These values correspond to the REDUND column of V\$ASM_TEMPLATE.

Table 7-5 Permitted Values for Oracle ASM Template Redundancy Attribute

Redundancy Attribute Value	Resulting Mirroring in Normal Redundancy Disk Group	Resulting Mirroring in High Redundancy Disk Group	Resulting Mirroring in External Redundancy Disk Group
MIRROR	Two-way mirroring	Three-way mirroring	(Not allowed)
HIGH	Three-way mirroring	Three-way mirroring	(Not allowed)
UNPROTECTED	No mirroring	(Not allowed)	No mirroring

Table 7-6 shows the initial attribute settings for the default templates. The type of mirroring associated with the **Mirroring** column for normal, high, and external redundancy disk groups is specified in Table 7-5. For example, the **Mirror** column for high redundancy disk groups displays MIRROR in Table 7-6. In Table 7-5, the MIRROR value for high redundancy disk groups corresponds to Three-way mirroring.

Table 7-6 Oracle ASM System Default Templates Attribute Settings

Template Name	File Striping	Mirroring, Normal Redundancy Disk Group	Mirroring, High Redundancy Disk Group	Mirroring, External Redundancy Disk Group	Primary Region	Mirror Region
CONTROLFILE	FINE	HIGH	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
DATAFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
ONLINELOG	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
ARCHIVELOG	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
TEMPFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
BACKUPSET	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
PARAMETERFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
DATAGUARDCONFIG	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
FLASHBACK	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
CHANGETRACKING	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
DUMPSET	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
XTRANSPORT	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
AUTOBACKUP	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
ASMPARAMETERFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
OCRFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD

[Table 7-7](#) shows the permitted primary and mirror region values for Oracle ASM templates. These values correspond to the PRIMARY_REGION and MIRROR_REGION columns of V\$ASM_TEMPLATE. For more information about disk regions, see "[Intelligent Data Placement](#)" on page 4-15.

Table 7-7 Permitted Values for Oracle ASM Template Primary and Mirror Region Attribute

Primary and Mirror Region Attribute Value	Description
COLD, MIRRORCOLD	Use the inner most tracks (closest to spindle) on the disk drive.
HOT, MIRRORHOT	Use the outermost tracks which have greater speed and higher bandwidth.

Adding Templates to a Disk Group

To add a template to a disk group, use the ADD TEMPLATE clause of the ALTER DISKGROUP statement.

If the name of your new template is not listed in [Table 7-6](#), then it is not used as a default template for database file types. To use the template, you must reference its name when creating a file.

When adding a template to a disk group, the attributes are optional. If no redundancy attribute is specified, the value defaults to MIRROR for a normal redundancy disk group, HIGH for a high redundancy disk group, and UNPROTECTED for an external redundancy disk group. If you do not specify a striping attribute, then the value defaults to COARSE. If you do not specify a value for the primary or mirror region attributes, then the value defaults to COLD and MIRRORCOLD respectively.

[Example 7-13](#) creates a new template named `reliable` for the normal redundancy disk group `data`.

Example 7-13 Adding a Template to a Disk Group

```
ALTER DISKGROUP data ADD TEMPLATE reliable ATTRIBUTES (HIGH FINE);
```

[Example 7-14](#) creates a new template named `unreliable` that specifies files are to be unprotected (no mirroring).

Example 7-14 Adding a Template to a Disk Group

```
ALTER DISKGROUP data ADD TEMPLATE unreliable ATTRIBUTES (UNPROTECTED);
```

Note: Oracle discourages using unprotected files unless you have implemented hardware mirroring. The previous example is presented only to further illustrate how the attributes for templates are set.

See Also: *Oracle Database SQL Language Reference* for more information about the ALTER DISKGROUP...ADD TEMPLATE command.

Modifying a Disk Group Template

The `MODIFY TEMPLATE` clause of the `ALTER DISKGROUP` statement enables you to modify the attribute specifications of an existing system default or user-defined disk group template. Only specified template attributes are changed. Unspecified properties retain their current value. When you modify an existing template, only new files created by the template reflect the attribute changes. Existing files maintain their attributes.

[Example 7-15](#) changes the striping attribute specification of the `reliable` template for disk group data.

Example 7-15 *Modifying a Disk Group Template*

```
ALTER DISKGROUP data MODIFY TEMPLATE reliable
    ATTRIBUTES (COARSE);
```

Dropping Templates from a Disk Group

Use the `DROP TEMPLATE` clause of the `ALTER DISKGROUP` statement to drop one or more templates from a disk group. You can only drop templates that are user-defined; you cannot drop system default templates.

[Example 7-16](#) drops the previously created template `unreliable` from data:

Example 7-16 *Dropping a Template from a Disk Group*

```
ALTER DISKGROUP data DROP TEMPLATE unreliable;
```

Creating Tablespaces in Oracle ASM: Specifying Attributes with Templates

Use the SQL `ALTER SYSTEM` and `CREATE TABLESPACE` statements to create a tablespace that uses a user-defined template to specify the attributes of the data file.

[Example 7-17](#) assumes that the template (`mytemplate`) has been defined.

Example 7-17 *Using a User-Defined Template to Specify Attributes*

```
ALTER SYSTEM SET DB_CREATE_FILE_DEST = '+data(mytemplate)';

CREATE TABLESPACE mytblspace;
```

Performing Oracle ASM Data Migration With RMAN

This chapter describes how to migrate data into and out of Oracle Automatic Storage Management (Oracle ASM) storage with Recovery Manager (RMAN).

This chapter includes the following topics:

- [Overview of Oracle ASM Data Migration](#)
- [Preparing to Migrate the Database to Oracle ASM Using RMAN](#)
- [Migrating the Database to Oracle ASM Using RMAN](#)
- [Migrating a Database from Oracle ASM to Alternative Storage](#)
- [Moving Data Files Between Oracle ASM Disk Groups Using RMAN](#)

See Also: *Oracle Database Backup and Recovery User's Guide* for complete information about using RMAN

Overview of Oracle ASM Data Migration

This section explains the basic concepts and tasks involved in migrating data to and from Oracle ASM.

This section includes the following topics:

- [Purpose of Oracle ASM Data Migration](#)
- [Basic Concepts of Oracle ASM Data Migration](#)
- [Basics Steps of Data Migration to Oracle ASM Using RMAN](#)

Purpose of Oracle ASM Data Migration

Alternatives to Oracle ASM storage include file systems, raw disks, and SAN configurations. Oracle ASM includes a number of benefits over these storage alternatives, including performance optimization, redundancy protection, and load balancing. You do not need to obtain a third-party Logical Volume Manager because Oracle ASM manages disks for you. Oracle Real Application Clusters (Oracle RAC) databases benefit from Oracle ASM because it provides ready-made shared storage.

If a database currently uses a storage system other than Oracle ASM, then you can migrate all or part of the database into Oracle ASM, thereby simplifying database administration. You can also migrate a fast recovery area to Oracle ASM.

Native operating system commands such as Linux `cp` or Windows `COPY` cannot write or read files in Oracle ASM storage. Because RMAN can read and write Oracle ASM

files, you can use RMAN to copy data files into and out of Oracle ASM storage or between Oracle ASM disk groups. This technique is useful if you must store backups on user-managed disks.

Basic Concepts of Oracle ASM Data Migration

You can migrate data to Oracle ASM with RMAN even if you are not using RMAN as your primary backup tool. The migration requires one RMAN database backup.

If you have sufficient disk space to hold the entire database both in Oracle ASM and alternative storage systems, then you can move a database directly into Oracle ASM. If you do not have sufficient storage, then you can back the database up to tape, create an Oracle ASM disk group that uses old disk space, and restore the database from tape to Oracle ASM.

After you set the location of the new recovery area, existing backups remain in the old recovery area and count against the total disk quota of the recovery area. The backups are deleted from the old recovery area when space is needed. These backups are usable by RMAN. You do not need to move legacy backups to the new Oracle ASM recovery area unless you need disk space. To free space consumed by files in the old recovery area, you can back them up to tape or migrate them to the Oracle ASM recovery area.

Note: A foreign archived redo log is a log received by a logical standby database for a LogMiner session. Foreign archived redo logs cannot be migrated. Unlike normal archived logs, foreign archived logs have a different DBID. For this reason, they cannot be backed up or restored on a logical standby database.

Migrating a database from Oracle ASM to an alternative storage system is similar to migration from an alternative storage system to Oracle ASM. The primary change is to modify each step to refer to file locations in the alternative storage system.

For information about migrating the database to Oracle ASM with Enterprise Manager, see [Chapter 9, "Administering Oracle ASM with Oracle Enterprise Manager"](#).

Basics Steps of Data Migration to Oracle ASM Using RMAN

This section discusses the process of migrating the entire database and fast recovery area from alternative storage to Oracle ASM using RMAN.

The fast recovery area is an optional disk location that you can use to store recovery-related files such as control file and online redo log copies, archived redo log files, flashback logs, and RMAN backups. Oracle Database and RMAN manage the files in the fast recovery area automatically. You can specify the disk quota, which is the user-specified maximum size of the fast recovery area. When the disk quota is reached, Oracle automatically deletes files that are no longer needed.

Flashback logs are Oracle-generated logs used to perform flashback database operations. The database can only write flashback logs to the fast recovery area. Flashback logs are written sequentially and are not archived. They cannot be backed up to disk.

To migrate the entire database and fast recovery area from alternative storage to Oracle ASM, perform the following steps:

1. Back up the database and server parameter file, and disable Oracle Flashback Database.

The Oracle Flashback Database option returns the entire database to a prior consistent SCN by means of the `FLASHBACK DATABASE` command in RMAN or SQL. A database flashback is different from traditional media recovery because it does not involve the restore of physical files, instead restoring your current data files to past states using saved images of changed data blocks. This feature uses flashback logs and archived redo logs.

This step is described in ["Preparing to Migrate the Database to Oracle ASM Using RMAN"](#) on page 8-3.

2. Restore files to Oracle ASM, recover the database, and optionally migrate the fast recovery area to Oracle ASM.

This step is described in ["Migrating the Database to Oracle ASM Using RMAN"](#) on page 8-5.

To migrate files from alternative storage to Oracle ASM, see ["Migrating a Database from Oracle ASM to Alternative Storage"](#) on page 8-9.

Preparing to Migrate the Database to Oracle ASM Using RMAN

This section explains how to prepare the database for migration. This section makes the following assumptions:

- You want to migrate the database to two Oracle ASM disk groups: `+DATA` for the database and `+FRA` for the fast recovery area.
- The database to be migrated to Oracle ASM storage is named `mydb`.

Note: If you do not want to migrate the fast recovery area, then skip step 10.

To prepare the database for Oracle ASM migration:

1. If the `COMPATIBLE` setting for the database is less than 11.0.0, then make any read-only transportable tablespaces read/write.

Read-only transportable tablespaces cannot be migrated because RMAN cannot back them up.

2. If the database is a physical standby database, and if managed recovery is started, then stop managed recovery.

A physical standby database is a copy of a production database that you can use for disaster protection.

For example, connect SQL*Plus to the database with `SYSDBA` privileges, and execute the following statement to stop managed recovery:

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE CANCEL;
```

Keep this terminal window open.

3. Copy the server parameter file or initialization parameter file to a temporary location.

The following example uses an operating system utility to copy the server parameter file:

```
$ cp spfileMYDB.ora orig_spfileMYDB.ora
```

4. In a new terminal window, start RMAN session and connect as `TARGET` to the database to be migrated. Optionally, connect to a recovery catalog.
5. Back up the data files to the Oracle ASM disk group.

The following example uses a `RUN` command to make a level 0 incremental backup and allocates four channels to increase the backup speed. A level 0 incremental backup is an RMAN incremental backup that backs up all data blocks in the data files being backed up. An incremental backup at level 0 is identical in content to a full backup, but unlike a full backup the level 0 backup is considered a part of the incremental backup strategy.

An incremental backup is an RMAN backup in which only modified blocks are backed up. Incremental backups are classified by **level**. A level 0 incremental backup performs the same function as a full backup in that they both back up all blocks that have ever been used. The difference is that a full backup does not affect blocks backed up by subsequent incremental backups, whereas an incremental backup affects blocks backed up by subsequent incremental backups.

A full backup is a non-incremental RMAN backup. Full does not refer to how much of the database is backed up, but to the fact that the backup is not incremental. Consequently, you can make a full backup of one data file.

Increase or decrease this number accordingly. The format clause specifies `+DATA`, which is the name of the Oracle ASM disk group to be used for storing the database.

```
RUN
{
  ALLOCATE CHANNEL dev1 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev2 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev3 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev4 DEVICE TYPE DISK;
  BACKUP AS COPY
    INCREMENTAL LEVEL 0
    DATABASE
    FORMAT '+DATA'
    TAG 'ORA_ASM_MIGRATION';
}
```

6. If block change tracking is enabled for the database, then optionally make a level 1 incremental backup that you can use later to recover the database copy.

Block change tracking is a database option that causes Oracle to track data file blocks affected by each database update. The tracking information is stored in a block change tracking file. When block change tracking is enabled, RMAN uses the record of changed blocks from the change tracking file to improve incremental backup performance by only reading those blocks known to have changed, instead of reading data files in their entirety.

The following example makes an incremental level 1 copy of the level 0 backup created in the previous step:

```
RUN
{
  ALLOCATE CHANNEL dev1 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev2 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev3 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev4 DEVICE TYPE DISK;
  BACKUP INCREMENTAL LEVEL 1
    FOR RECOVER OF COPY WITH TAG 'ORA_ASM_MIGRATION'
    DATABASE;
```

}

7. If the database is in ARCHIVELOG mode, and if the database is open, then archive the online logs.

The following example uses the SQL command to archive the current redo logs:

```
RMAN> SQL "ALTER SYSTEM ARCHIVE LOG CURRENT";
```

8. If the database instance is currently using a server parameter file, then back it up.

The following example backs up the server parameter file:

```
RMAN> BACKUP AS BACKUPSET SPFILE;
```

9. If block change tracking is enabled, then disable it.

The following command disables block change tracking:

```
RMAN> SQL "ALTER DATABASE DISABLE BLOCK CHANGE TRACKING";
```

10. If Flashback Database is enabled, then disable it and drop any guaranteed restore points.

Note: If you are not migrating the recovery area, then skip this step.

Disabling Oracle Flashback Database is necessary because you cannot migrate flashback logs to Oracle ASM. The following command disables Flashback Database:

```
RMAN> SQL "ALTER DATABASE FLASHBACK OFF";
```

The following command drops the guaranteed restore point named Q106:

```
RMAN> SQL "DROP RESTORE POINT Q106";
```

11. Shut down the database consistently.

The following command shuts down the database:

```
RMAN> SHUTDOWN IMMEDIATE;
```

Migrating the Database to Oracle ASM Using RMAN

The following procedure is intended to minimize database downtime. The steps differ slightly depending on whether you are migrating a primary or standby database. The procedure makes the same assumptions described in ["Preparing to Migrate the Database to Oracle ASM Using RMAN"](#) on page 8-3. If you are not migrating the recovery area to Oracle ASM, then you must modify some of the steps, which are noted.

Note: The following procedure switches between SQL*Plus and RMAN, so keep a terminal window open for each utility.

To migrate the database to Oracle ASM:

1. Follow the steps in ["Preparing to Migrate the Database to Oracle ASM Using RMAN"](#) on page 8-3.
2. Restore or create a server parameter file in Oracle ASM storage.

The steps depend on whether the database is using a server parameter file:

- If the database is using a server parameter file, then restore it to the Oracle ASM disk group with the following commands, where *sid* is the SID of the instance:

```

RMAN> STARTUP MOUNT;
RMAN> RESTORE SPFILE TO '+DATA/spfilesid.ora';
RMAN> SHUTDOWN IMMEDIATE;

```

- If the database is not using a server parameter file, then create one in Oracle ASM. Execute the `CREATE SPFILE` command in SQL*Plus as follows, where *sid* is the SID of the database (the command spans two lines):

```
SQL> CREATE SPFILE='+DATA/spfilesid.ora' FROM PFILE='?/dbs/initssid.ora';
```

Afterward, delete *spfilesid.ora* and *initssid.ora* from the *?/dbs* directory and create a new *initssid.ora* with the following line of content:

```
SPFILE='+DATA/spfilesid.ora'
```

3. Set Oracle Managed Files initialization parameters to Oracle ASM locations.

Note: If you are not migrating the fast recovery area, then do not change the `DB_RECOVERY_FILE_DEST` and `DB_RECOVERY_FILE_DEST_SIZE` initialization parameter settings. However, you must set `DB_CREATE_ONLINE_LOG_DEST_n` parameter to an Oracle ASM location for migration of the online redo logs.

Set the `DB_CREATE_FILE_DEST` and optional `DB_CREATE_ONLINE_LOG_DEST_n` initialization parameters to Oracle ASM disk groups. If the database uses a recovery area, then change the recovery area location to the Oracle ASM disk group. Also, change the recovery area size.

Execute commands in SQL*Plus as shown in the following example. The example assumes that the size of the fast recovery area is 100 GB and specifies the disk group `+FRA` for the fast recovery area.

```

SQL> STARTUP FORCE NOMOUNT;
SQL> ALTER SYSTEM SET DB_CREATE_FILE_DEST='+DATA' SID='*';
SQL> ALTER SYSTEM SET DB_RECOVERY_FILE_DEST_SIZE=100G SID='*';
SQL> ALTER SYSTEM SET DB_RECOVERY_FILE_DEST='+FRA' SID='*';

```

4. Set the `CONTROL_FILES` initialization parameter to Oracle ASM locations.

If you *are* migrating the fast recovery area, then enter the following commands in SQL*Plus to restart the database instance and set the control file locations to disk groups `+DATA` and `+FRA`:

```

SQL> STARTUP FORCE NOMOUNT;
SQL> ALTER SYSTEM SET CONTROL_FILES='+DATA','+FRA' SCOPE=SPFILE SID='*';

```

If you are *not* migrating the fast recovery area, then enter the following commands in SQL*Plus to restart the database instance and set the control file locations to disk group `+DATA`:

```

SQL> STARTUP FORCE NOMOUNT;
SQL> ALTER SYSTEM SET CONTROL_FILES='+DATA','+DATA' SCOPE=SPFILE SID='*';

```

5. Migrate the control file to Oracle ASM and mount the control file.

Switch to the RMAN terminal to restore the control file. In the following example, *original_cf_name* is a control file name in the initialization parameter file before migration:

```
RMAN> STARTUP FORCE NOMOUNT;
RMAN> RESTORE CONTROLFILE FROM 'original_cf_name';
RMAN> ALTER DATABASE MOUNT;
```

6. Migrate the data files to Oracle ASM.

Use RMAN to switch to the database copy that you created in step 5 in "[Preparing to Migrate the Database to Oracle ASM Using RMAN](#)" on page 8-3. The switch renames all the data files to files on Oracle ASM disk groups. Afterward, recover the database. If incremental backups were taken, then RMAN applies them during recovery. For example, enter the following commands at the RMAN prompt:

```
SWITCH DATABASE TO COPY;
RUN
{
  ALLOCATE CHANNEL dev1 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev2 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev3 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev4 DEVICE TYPE DISK;
  RECOVER DATABASE;
}
```

7. If the database uses block change tracking or Flashback Database, then enable these features.

Note: If you are not migrating the recovery area, then you do not need to enable Flashback Database because you did not disable it.

For example, enter the following statements in SQL*Plus:

```
SQL> ALTER DATABASE ENABLE BLOCK CHANGE TRACKING USING FILE '+DATA';
SQL> ALTER DATABASE FLASHBACK ON;
```

8. Place the database in its normal operation mode.

The normal operational mode depends on whether the database is a primary or standby database:

- If the database is a primary database, then open it as follows:

```
SQL> ALTER DATABASE OPEN;
```

- If the database is a standby database, then resume managed recovery mode as follows:

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE;
```

9. Drop the tempfiles and re-create them in Oracle ASM.

Use SQL*Plus to re-create the tempfiles. In the following example, the name of the tempfile in the original storage is *tempfile_name*. The name of the temporary tablespace is *temp_tbs_name*.

```
SQL> ALTER DATABASE TEMPFILE 'tempfile_name' DROP;
SQL> ALTER TABLESPACE temp_tbs_name ADD TEMPFILE;
```

10. Migrate the online redo log files.

If this is a primary database, then add new log group members in Oracle ASM and drop the old members. You can use the following PL/SQL script to migrate the online redo log groups into an Oracle ASM disk group. The PL/SQL script assumes that the Oracle Managed Files initialization parameters specified in step 3 are set.

Example 8–1 Migrating the Online Redo Logs

```

SET SERVEROUTPUT ON;
DECLARE
  CURSOR rlc IS
    SELECT GROUP# GRP, THREAD# THR, BYTES, 'NO' SRL
    FROM   V$LOG
    UNION
    SELECT GROUP# GRP, THREAD# THR, BYTES, 'YES' SRL
    FROM   V$STANDBY_LOG
    ORDER BY 1;
  stmt    VARCHAR2(2048);
BEGIN
  FOR rlcRec IN rlc LOOP
    IF (rlcRec.srl = 'YES') THEN
      stmt := 'ALTER DATABASE ADD STANDBY LOGFILE THREAD ' ||
              rlcRec.thr || ' SIZE ' || rlcRec.bytes;
      EXECUTE IMMEDIATE stmt;
      stmt := 'ALTER DATABASE DROP STANDBY LOGFILE GROUP ' || rlcRec.grp;
      EXECUTE IMMEDIATE stmt;
    ELSE
      stmt := 'ALTER DATABASE ADD LOGFILE THREAD ' ||
              rlcRec.thr || ' SIZE ' || rlcRec.bytes;
      EXECUTE IMMEDIATE stmt;
      BEGIN
        stmt := 'ALTER DATABASE DROP LOGFILE GROUP ' || rlcRec.grp;
        DBMS_OUTPUT.PUT_LINE(stmt);
        EXECUTE IMMEDIATE stmt;
      EXCEPTION
        WHEN OTHERS THEN
          EXECUTE IMMEDIATE 'ALTER SYSTEM SWITCH LOGFILE';
          EXECUTE IMMEDIATE 'ALTER SYSTEM CHECKPOINT GLOBAL';
          EXECUTE IMMEDIATE stmt;
      END;
    END IF;
  END LOOP;
END;
/

```

11. Optionally, migrate backups and copies in the old fast recovery area to Oracle ASM as follows:

- a.** If foreign archived logs exists in the recovery area, then you cannot migrate them to Oracle ASM. Execute the following command at the RMAN prompt:

```
RMAN> DELETE REMOTE ARCHIVELOG ALL;
```

- b.** Back up archived redo log files, backup sets, and data file copies to Oracle ASM. For example, execute the following command at the RMAN prompt:

```

RUN
{
  ALLOCATE CHANNEL dev1 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev2 DEVICE TYPE DISK;
  ALLOCATE CHANNEL dev3 DEVICE TYPE DISK;

```

```

ALLOCATE CHANNEL dev4 DEVICE TYPE DISK;

BACKUP AS COPY ARCHIVELOG ALL DELETE INPUT;
BACKUP BACKUPSET ALL DELETE INPUT;
BACKUP AS COPY DATAFILECOPY ALL DELETE INPUT;
}

```

Migrating a Database from Oracle ASM to Alternative Storage

Migrating a database from Oracle ASM to an alternative storage system is essentially the reverse of the migration to Oracle ASM. Modify the steps in ["Preparing to Migrate the Database to Oracle ASM Using RMAN"](#) on page 8-3 and ["Migrating the Database to Oracle ASM Using RMAN"](#) on page 8-5 as follows:

- If the procedure specifies Oracle Managed Files locations, then alter the procedure to use locations in alternative storage.
- If the `FORMAT` clause of the `BACKUP` command specifies an Oracle ASM location, then change the backup format to an alternative storage location.
- If a filename used in a SQL statement is an Oracle ASM location, then change it to a filename in the alternative storage location.

Moving Data Files Between Oracle ASM Disk Groups Using RMAN

You may want to move an active data file in an `ARCHIVELOG` mode database from one Oracle ASM disk group to another. In this case, you use `BACKUP AS COPY` to copy the data file to the new disk group and `SET NEWNAME` and `SWITCH` commands to rename the data file in the control file.

For this scenario, assume that you are using disk groups `DATA` and `USERDATA`. You want to move data file `+DATA/orcl/datafile/users.261.689589837` to disk group `USERDATA`.

Ensure that `ARCHIVELOG` mode is enabled for the database before beginning the procedure to move data files.

To move a data file from one Oracle ASM disk group to another disk group, perform the following steps.

1. Start RMAN and connect to the target database.

For example:

```

$ rman
RMAN> CONNECT TARGET SYS@orcl
target database Password: XXXXXXXXX
connected to target database: ORCL (DBID=1217369048)

```

2. Generate a report that shows the names of the data files.

Execute the following `REPORT` command after connecting RMAN to the target database. Note the data file name of the file to be moved.

For example:

```
RMAN> REPORT SCHEMA;
```

```
Report of database schema for database with db_unique_name ORCL
```

```
List of Permanent Datafiles
=====
```

File	Size(MB)	Tablespace	RB segs	Datafile Name
1	740	SYSTEM	***	+DATA/orcl/datafile/system.258.689589737
2	570	SYSAUX	***	+DATA/orcl/datafile/sysaux.259.689589785
3	55	UNDOTBS1	***	+DATA/orcl/datafile/undotbs1.260.689589831
4	5	USERS	***	+DATA/orcl/datafile/users.261.689589837

List of Temporary Files

File	Size(MB)	Tablespace	Maxsize(MB)	Tempfile Name
1	20	TEMP	32767	+DATA/orcl/tempfile/temp.262.689589851

3. Back up the data file to the new Oracle ASM disk group.

Issue the `BACKUP AS COPY` command to back up the data file on DATA to USERDATA.

For example:

```
RMAN> BACKUP AS COPY
      DATAFILE "+DATA/orcl/datafile/users.261.689589837"
      FORMAT "+USERDATA";
```

```
Starting backup at 16-JUN-09
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=51 device type=DISK
channel ORA_DISK_1: starting datafile copy
input datafile file number=00004 name="+DATA/orcl/datafile/users.261.689589837"
output file name="+USERDATA/orcl/datafile/users.256.689682663"
tag=TAG20090616T103101 RECID=13 STAMP=689682663
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:01
Finished backup at 16-JUN-09
```

You could also specify the data file by the data file number and data file type.

For example:

```
BACKUP AS COPY
      DATAFILE 4
      FORMAT "+USERDATA";
```

4. Offline the data file that you intend to move to a new disk group.

Execute the following `SQL` command in the RMAN client. Note that you use two single quotation marks around the name of the data file, not double quotation marks.

For example:

```
RMAN> SQL "ALTER DATABASE DATAFILE
          '+'DATA/orcl/datafile/users.261.689589837'" OFFLINE";
```

```
sql statement: ALTER DATABASE DATAFILE
          '+'DATA/orcl/datafile/users.261.689589837'" OFFLINE
```

5. Point the control file to the newly created copy of the data file.

Run the `SWITCH . . . TO COPY` command in the RMAN client. The `TO COPY` option of `SWITCH` switches the data file to the most recent copy of the data file.

For example:

```
RMAN> SWITCH DATAFILE "+DATA/orcl/datafile/users.261.689589837" TO COPY;
```

```
datafile 4 switched to datafile copy
"+USERDATA/orcl/datafile/users.256.689682663"
```

The output of this command displays the new name of the data file.

6. Recover the renamed data file.

Run the RECOVER command in the RMAN client.

For example:

```
RMAN> RECOVER DATAFILE "+USERDATA/orcl/datafile/users.256.689682663";

Starting recover at 16-JUN-09
using channel ORA_DISK_1
starting media recovery
media recovery complete, elapsed time: 00:00:01
Finished recover at 16-JUN-09
```

7. Bring the data file online.

Execute the SQL command in the RMAN client. Note that you use two single quotation marks around the name of the data file, not double quotation marks.

For example:

```
RMAN> SQL "ALTER DATABASE DATAFILE
'+USERDATA/orcl/datafile/users.256.689682663' ONLINE";

sql statement: ALTER DATABASE DATAFILE
'+USERDATA/orcl/datafile/users.256.689682663' ONLINE
```

8. Delete the data file copy from the original Oracle ASM disk group.

In this scenario, +DATA/orcl/datafile/users.261.689589837 is the original data file in DATA. Because you issued SET NEWNAME and SWITCH commands for this data file, the original file is now recorded in the RMAN repository as a data file copy. Execute a DELETE command in the RMAN client to remove this file.

For example:

```
RMAN> DELETE DATAFILECOPY "+DATA/orcl/datafile/users.261.689589837";

released channel: ORA_DISK_1
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=51 device type=DISK
List of Datafile Copies
=====

Key      File S Completion Time Ckp SCN    Ckp Time
-----  - - - - -
14       4    A 16-JUN-09      864471    16-JUN-09
Name: +DATA/orcl/datafile/users.261.689589837
Tag: TAG20090615T084217
```

```
Do you really want to delete the above objects (enter YES or NO)? y
deleted datafile copy
datafile copy file name=+DATA/orcl/datafile/users.261.689589837 RECID=14
STAMP=689683255
Deleted 1 objects
```

Administering Oracle ASM with Oracle Enterprise Manager

This chapter describes how to administer Oracle Automatic Storage Management (Oracle ASM) by using Oracle Enterprise Manager to provision and manage data file storage. This chapter contains the following topics:

- [Oracle Automatic Storage Management Home Page](#)
- [Configuring Oracle ASM Initialization Parameters with Oracle Enterprise Manager](#)
- [Managing Oracle ASM Users with Oracle Enterprise Manager](#)
- [Managing Disk Groups with Oracle Enterprise Manager](#)
- [Managing Oracle ASM File Access Control with Oracle Enterprise Manager](#)
- [Managing Directories, Files, and Aliases with Oracle Enterprise Manager](#)
- [Managing Disk Group Templates with Oracle Enterprise Manager](#)
- [Monitoring Oracle ASM Performance with Oracle Enterprise Manager](#)
- [Backing Up Oracle ASM Files with Oracle Enterprise Manager](#)
- [Performing Bad Block Recovery with Oracle Enterprise Manager](#)
- [Migrating to Oracle ASM with Oracle Enterprise Manager](#)
- [Oracle ASM Support Workbench](#)
- [Oracle ASM: Oracle by Example Series](#)


Oracle Automatic Storage Management Home Page

All Oracle ASM administration tasks begin with the Oracle Automatic Storage Management Home page in Oracle Enterprise Manager, shown in [Figure 9-1](#). The Oracle Automatic Storage Management Home page displays:

- The status of the Oracle ASM instance.
- A chart that shows the used and free space of each disk group and disk group internal usage.
- A list of databases that are serviced by the Oracle ASM instance.
- A list of Oracle Automatic Storage Management Cluster File Systems that are serviced by the Oracle ASM instance.
- A list of other non-Oracle ACFS volumes.

- A list of alerts for the Oracle ASM instance and the host computer.
- Links to the Oracle ASM Performance, Disk Groups, Configuration, Users, and Oracle ACFS pages.

Figure 9–1 Oracle Automatic Storage Management Home Page



Accessing the Oracle ASM Home Page in Single-Instance Oracle Databases

To access the Oracle Automatic Storage Management Home page on a single-instance system:

1. Access the Database Home page by logging in as any database user.

See Also: *Oracle Database 2 Day DBA* for more information

2. Under the heading labeled **General** on the left side of the Database Instance Home page as shown in Figure 9–1, click the link next to the last entry in the General column labeled **Oracle ASM**.
3. If prompted for Oracle ASM login credentials, then enter the user `SYS`, provide the `SYS` password that was set for the Oracle ASM instance upon installation, and connect as `SYSASM`. Oracle Enterprise Manager displays the Oracle Automatic Storage Management Home page similar to the page in Figure 9–1.

See Also: "Authentication for Accessing Oracle ASM Instances" on page 3-22 for more information about authentication

Accessing the Oracle ASM Home Page in Oracle RAC Databases

To access the Oracle Automatic Storage Management Home page in an Oracle Real Application Clusters (RAC) environment:

1. Log in to Oracle Enterprise Manager Database Control on any node that is running the Oracle Management Service (OMS).

OMS is automatically started on the node on which Database Configuration Assistant (DBCA) was originally run to create the cluster database. Depending on your configuration, OMS might also be running on other nodes.

See Also: *Oracle Database 2 Day + Real Application Clusters Guide* for more information

2. On the Cluster Database page, under the Instances heading, click the link for the Oracle ASM instance that you want to access.

Configuring Oracle ASM Initialization Parameters with Oracle Enterprise Manager

You can configure Oracle ASM with the Oracle Enterprise Manager Configuration Parameters page similar to the one shown in [Figure 9–2](#).

Figure 9–2 Oracle Automatic Storage Management Configuration Parameters Page

The screenshot shows the 'Configuration Parameters' page for Oracle ASM. The page has a navigation bar with tabs: Home, Performance, Disk Groups, Configuration (selected), Users, and ASM Cluster File System. Below the navigation bar, the 'Configuration Parameters' section is displayed. It includes the following parameters:

- Disk Discovery Path:** A text input field containing '/devices/disk*'. Below it is a description: 'This path limits the set of disks considered for discovery. It should match the path or the directory containing the disk. e.g. /dev/raw/* for Linux based operating systems.' To the right are 'Revert' and 'Apply' buttons.
- Auto Mount Disk Groups:** A text input field containing 'DATA, FRA'. Below it is a description: 'The list of the Disk Group names to be mounted by the ASM at startup or when ALTER DISKGROUP ALL MOUNT command is used.'
- Rebalance Power:** A dropdown menu set to '1'. Below it is a description: 'Higher values allows the operation to complete more quickly but takes more I/O bandwidth away from the database. Lower values causes rebalance to take longer but leave more I/O bandwidth for the database.'
- Preferred Read Failure Groups:** An empty text input field. Below it is a description: 'Specify a comma-separated list of failure groups whose member disks will be preferred read disks for this node. If there is more than one mirror copy to read from, ASM will read from the preferred disk.'

To configure the Oracle ASM instance:

1. Access the Oracle Automatic Storage Management Home page.
See "[Oracle Automatic Storage Management Home Page](#)" on page 9-1 for instructions.
2. Click the **Configuration** link at the top of the page to view the Configuration Parameters page.
3. Update the configuration options on the page:
 - Disk Discovery Path (ASM_DISKSTRING initialization parameter)
Enter a string value. See "[ASM_DISKSTRING](#)" on page 3-7.
 - Auto Mount Disk Groups (ASM_DISKGROUPS initialization parameter)


- Enter a string value. See "[ASM_DISKGROUPS](#)" on page 3-7.
 - Rebalance Power (ASM_POWER_LIMIT initialization parameter)
Enter an integer value. See "[ASM_POWER_LIMIT](#)" on page 3-8.
 - Preferred Read Failure Groups (Only in Oracle RAC environments) (ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter)
Enter a comma-delimited string of values in which each string represents a failure group. This parameter is generally used only for clustered Oracle ASM instances. See "[ASM_PREFERRED_READ_FAILURE_GROUPS](#)" on page 3-8.
4. Click **Apply** to save your changes or click **Revert** to discard any changes made on the page.

Bringing Disks Online and Offline

Perform the following procedures to bring a disk online or offline:

1. Access the Oracle Automatic Storage Management Home page.
See "[Oracle Automatic Storage Management Home Page](#)" on page 9-1 for instructions.
2. Click the **Disk Groups** link at the top of the page to view the Disk Groups page.
Oracle Enterprise Manager displays a page similar to the one show in [Figure 9-3](#).

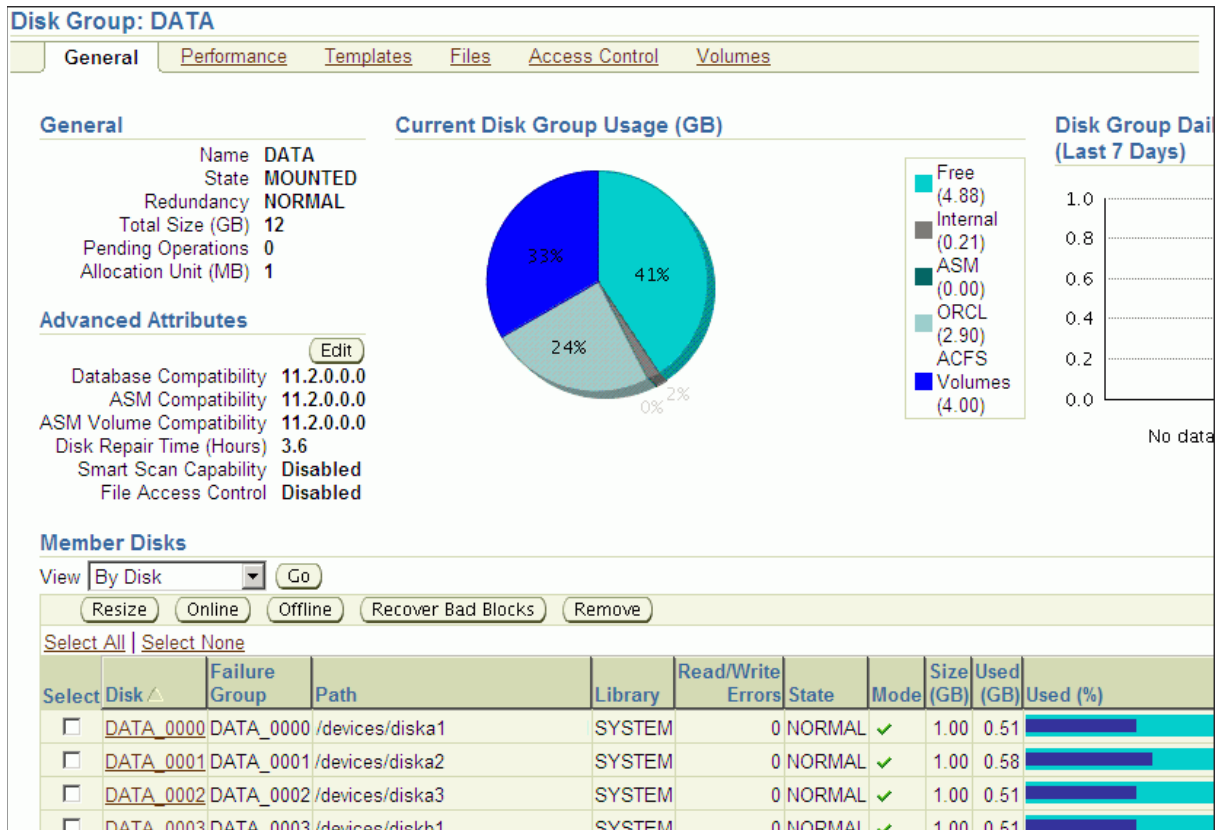
Figure 9-3 Oracle Automatic Storage Management Disk Groups List Page



Select	Name	State	Redundancy	Size (GB)	Used (GB)	Used (%)	Usable Free (GB)	Member Disks
<input type="checkbox"/>	DATA	MOUNTED	NORMAL	12.00	7.12	59.33	1.89	12
<input type="checkbox"/>	FRA	MOUNTED	NORMAL	12.00	0.70	5.85	5.10	12

TIP The usable free space specifies the amount of space that can be safely used for data. A value above zero means that redundancy can be properly restored after a disk failure.
 TIP Mount All and Dismount All operation will only mount and dismount the disk groups specified in the Auto Mount Disk Groups parameter.

3. Click the name of a disk in the Name column and Oracle Enterprise Manager displays the General page for the disk group that you want to modify. The General page that you see should be similar to the one shown in [Figure 9-4](#).

Figure 9–4 Oracle Automatic Storage Management Disk Group General Page


4. To bring a disk online, select the disk that you want to bring online and click **Online**. If you bring a disk online, then Oracle ASM performs the following:
 - Oracle ASM performs the online operation on the selected disks. You can select one or more or disks, or all of the disks to bring online. If you select all of the disks, then Oracle uses the `ALL` keyword for this operation. Otherwise, only the selected disks are affected by the online operation.
 - During the online process, Oracle Enterprise Manager first places a disk into online mode to accommodate write operations. Then Oracle Enterprise Manager resynchronizes stale data on the disk with the current data until the disk is fully online for read operations; this is the Oracle ASM fast mirror resync feature.
 - You can also choose the By Failure Group view to perform this operation on a selected disk group. In this case, Oracle uses the `SQL` for disks in failure group.
 - You can also select a disk that is online, but Oracle Enterprise Manager ignores that selection and only brings disks online that were previously offline.
 - Click **Yes** on the Oracle Enterprise Manager Confirmation dialog to begin bringing the disk online or click **No** to cancel the operation. You can also click **Show SQL** to review the SQL that Oracle Enterprise Manager uses for the online operation.
5. To take a disk offline, select the disk that you want to take offline and click **Offline**. If you bring a disk offline, then Oracle ASM performs the following:

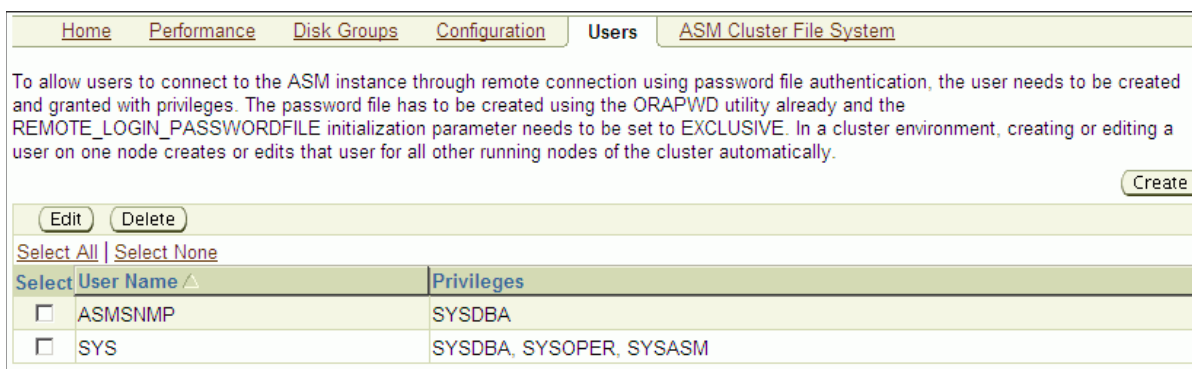
- Depending on how much data is stale, the operation may require additional time to complete. You can monitor the operation's progress using the Pending Operations link from the Disk Group Home page.
- The value that you set for Rebalance Power using `ASM_POWER_LIMIT` must be a whole number ranging from 0 to 11, inclusively. The text on the page describes how the value of Rebalance Power affects the fast mirror resync operation. If the `ASM_POWER_LIMIT` is set to 0, then Oracle ASM does not perform automatic rebalance.
- By default, the Rebalance Power value uses the value of the `ASM_POWER_LIMIT` initialization parameter.
- Oracle Enterprise Manager displays an Offline Confirmation page and the text on this page explains that after Oracle Enterprise Manager takes the selected disks offline, the Oracle ASM instance waits for the period of time that is specified by the value that you have set for the `DISK_REPAIR_TIME` initialization parameter before Oracle ASM drops the disks. During the disk repair time period, you can repair the disks and then place them back online. For each disk that you successfully place online, the pending drop operation is canceled and a resynchronization operation begins.
- If you have set `DISK_REPAIR_TIME` to a large value, for example, greater than one day, then this increases the possibility of data loss. This is because the files in the disk group are protected with a lower degree of redundancy during this period.

Click **Yes** on the Oracle Enterprise Manager Confirmation dialog to begin the offline processing or click **No** to stop the operation. You can also click **Show SQL** to see the SQL that Oracle Enterprise Manager uses for the offline operation.

Managing Oracle ASM Users with Oracle Enterprise Manager

You can manage Oracle ASM users with the Oracle Enterprise Manager Users page, shown in [Figure 9-5](#). Access this page by clicking the Users tab that is near the top center of the Oracle Automatic Storage Management Home page.

Figure 9-5 Oracle Automatic Storage Management Users Page



To manage Oracle ASM users:

1. Access the Oracle Automatic Storage Management Home page.

See ["Oracle Automatic Storage Management Home Page"](#) on page 9-1 for instructions.

2. Click the **Users** link at the top to view the Users property page.
3. Click **Create** to add a new Oracle ASM user. When you add users, Oracle Enterprise Manager displays a Create User page in which you must enter user login credentials that include the user name, password, and password confirmation. You can also grant privileges to the new user by selecting privileges in the Available Privileges column and clicking the right arrow buttons to move privileges to the Granted Privileges column, or move the privilege by double clicking it. The available privileges include SYSDBA, SYSOPER, and SYSASM. Click **OK** to create the user, **Cancel** to cancel the procedure, or **Show SQL** to view the SQL that Oracle Enterprise Manager uses to create the user.
4. To edit a user's properties, click the box in the Select column for the user that you want to edit and click **Edit**. On the Edit User page, modify the password if needed. You can also alter the privileges that are assigned to the user by selecting the privilege and using the right and left arrow buttons to move the privilege from the Granted Privileges column to the Available Privileges column and vice versa. You can also move the privilege by double clicking it. Click **OK** to edit the user properties, **Revert** to cancel the procedure, or **Show SQL** to view the SQL that Oracle Enterprise Manager uses to edit the user's properties.
5. To delete an Oracle ASM user, click the box in the Select column for the user that you want to delete and click **Delete**. On the confirmation dialog, click **Yes** to delete the user or **No** to stop the user deletion procedure.

See Also: ["Authentication for Accessing Oracle ASM Instances"](#) on page 3-22 for more information about user authentication

Managing Disk Groups with Oracle Enterprise Manager

This section describes how to manage disk groups under the following topics:

- [Creating Disk Groups](#)
- [Adding Disks to Disk Groups](#)
- [Dropping Disks from Disk Groups](#)
- [Dropping Disk Groups](#)
- [Monitoring Disk Group Usage](#)
- [Mounting and Dismounting Disk Groups](#)
- [Administering Advanced Disk Group Properties](#)

Creating Disk Groups

You can create additional Oracle ASM disk groups to:

- Have disk groups with different redundancy levels (normal, high, or external), depending on availability requirements and storage system capabilities.
- Separate different classes of storage, such as SCSI drives and SATA drives, into different disk groups. Disks in a disk group should have similar size and performance characteristics.
- Store the fast recovery area in a separate disk group from the database.

To create a disk group:

1. Access the Oracle Automatic Storage Management Home page.
See "Oracle Automatic Storage Management Home Page" on page 9-1 for instructions.
2. Click the **Disk Groups** link at the top of the page to display the Disk Groups property page.
3. If you are not logged in, then the Log In page appears and you should log in as the SYS user, connecting as SYSASM. Provide the SYS password for the Oracle ASM instance. This password was set up when the Oracle ASM instance was created.
4. Click **Create** and Oracle Enterprise Manager displays a Create Disk Group page similar to the page in Figure 9-6. The Create Disk Group page displays a list of Oracle ASM disks that are available to be added to a disk group. This includes disks with a header status of CANDIDATE, PROVISIONED, or FORMER, and so on.

Figure 9-6 Oracle Automatic Storage Management Create Disk Group Page

Create Disk Group Show SQL Cancel OK

* Name

Redundancy HIGH NORMAL EXTERNAL

Allocation Unit (MB)

An allocation unit (AU) is the fundamental unit in which contiguous disk space is allocated to ASM files. ASM file extent size is a multiple of AUs. The AU size cannot be modified later.

Candidate Member Disks

Select	Path	Header Status	Library	Label	ASM Disk Name	Size	Unit	Force Reuse	Failure Group
<input type="checkbox"/>	/devices/diske1	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diske2	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diske3	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diskk1	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diskk2	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	

5. If you want to display not only disks that are available to be added to a disk group, but all Oracle ASM disks, including disks that belong to a disk group (header status = MEMBER), then select **All Disks** from the Select Member Disks list.
The page is re-displayed with the complete list of Oracle ASM disks that were discovered by the Oracle ASM instance.
6. Enter the disk name, select a redundancy type (high, normal, or external), and optionally, you can enter a failure group name.
7. Select disks as follows:
 - Select the box to the left of each disk that you want to include in the new disk group.
 - Optionally, you can enter a disk name, the size of the disk that you want Oracle ASM to use, and the failure group name.
 - If you want to force the inclusion of a disk in the disk group, select the **Force Usage** box for that disk.

Caution: The Force Usage box causes the disk to be added to the new disk group even if the disk belongs to another disk group and has valid database data. This data will be lost. You must be certain that you are selecting a disk that can legitimately be added to the disk group. See ["Creating Disk Groups"](#) on page 4-3 for a discussion of the FORCE option.

- Specify a **Disk Compatibility** value for the disk group to enable your system to use the latest Oracle ASM features. See ["Configuring Disk Group Compatibility Attributes"](#) on page 9-15.

Note: You cannot reverse setting a disk compatibility value.

For information about disk group compatibility attributes, see ["Disk Group Compatibility"](#) on page 4-31.

- Click **OK** to create the disk group or click **Cancel** to stop the disk group creation operation. You can also click **Show SQL** to show the SQL statements that the disk group creation operation uses.

See Also: See ["Creating Disk Groups"](#) on page 4-3 for more information about creating disk groups the *Oracle Database Reference* description of the V\$ASM_DISK view for information about the various header statuses of Oracle ASM disks

Adding Disks to Disk Groups

You can add disks to a disk group to increase the total amount of storage space in a disk group. If you add multiple disks in a single operation, then Oracle ASM rebalances the disk group to evenly distribute the data on all disks, including the newly added disks.

You can control the power of a rebalance operation, which is a number from 0 to 11. The higher the number, the faster the rebalance operation completes. Lower numbers cause rebalancing to take longer, but consume fewer processing and I/O resources. This leaves these resources available for the database. The default value of 1 minimizes disruption to the database.

A value of 0 blocks the rebalance operation altogether. Manual or automatic rebalance can then occur at a later time. You might want to postpone rebalancing to wait for a time when there are fewer demands on the database, or to add more disks or drop disks later if you want the rebalancing to be done only once for all of your disk group changes.

See Also: ["Altering Disk Groups"](#) on page 4-11 for information about controlling rebalance operations

To add one or more disks to a disk group:

- Access the Oracle Automatic Storage Management Home page.

See Also: ["Oracle Automatic Storage Management Home Page"](#) on page 9-1 for instructions

- Click the **Disk Groups** link to display the Disk Groups property page.

3. If the Oracle ASM Login page appears, then log in as the *SYS* user, connecting as *SYSASM*. Provide the *SYS* password for the Oracle ASM instance. This password was set up when the Oracle ASM instance was created.
4. Click a link in the **Name** column to select the disk group to which you want to add disks. The Disk Group page similar to the page in [Figure 9-4](#) on page 9-5 appears, displaying a list of disks that are in the disk group.
5. Click **Add** and Oracle Enterprise Manager displays the Add Disks page. This page displays a list of Oracle ASM disks that are available to be added to the disk group. This includes disks with the header status of *CANDIDATE*, *PROVISIONED*, or *FORMER*, and so on.

Figure 9-7 Oracle Automatic Storage Management Add Disks Page

Add Disks Show SQL Cancel OK

Rebalance Power
 Rebalance operations redistribute data evenly across all drives. ASM automatically rebalances a disk group whenever disks are added or dropped. To manually rebalance all disk groups, you must specify the Rebalance Power. Higher values use more I/O bandwidth and complete rebalance more quickly. Lower values cause rebalance to take longer, but use less I/O bandwidth. Values range from 0 to 11.
 Rebalance Power

Candidate Member Disks

Select Member Disks Only Candidate D

Select	Path	Header Status	Library	Label	ASM Disk Name	Size	Unit	Force Reuse	Failure Group
<input type="checkbox"/>	/devices/diske1	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diske2	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diske3	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	

6. To display not only disks that are available to be added to a disk group, but all of the Oracle ASM disks, including disks that belong to a disk group (header status = *MEMBER*), select **All Disks** from the Select Member Disks drop-down list on the right hand side of the page. Oracle Enterprise Manager re-displays the page with the complete list of Oracle ASM disks that were discovered by the Oracle ASM instance.
7. Optionally change the rebalance power by selecting a value from the Rebalance Power list.
8. Select disks as follows:
 - Select the box to the left of each disk that you want to add to the disk group.
 - If you want to force the inclusion of a disk in the disk group, select the **Force Reuse** box in the right-most column.

Caution: The Force Reuse box causes the disk to be added to the disk group even if the disk belongs to another disk group and has valid database data. This data will be lost. You must be certain that you are selecting a disk that can legitimately be added to the disk group. Refer to ["Adding Disks to a Disk Group"](#) on page 4-12 for an example of the **FORCE** option.

- Optionally enter an Oracle ASM disk name for each disk. Oracle ASM provides a name if you do not.
9. Click **OK** to add the selected disks.

See Also: ["Altering Disk Groups"](#) on page 4-11 for more information about modifying disk group properties

Dropping Disks from Disk Groups

When you drop a disk from a disk group, Oracle ASM rebalances the disk group by moving all of the file extents from the dropped disk to other disks in the disk group. Oracle ASM then releases the disk and you can then add the disk to another disk group or use it for other purposes. You can drop one or more disks in a single operation. You can also optionally set rebalance power for the drop operation. The following is a list of possible reasons for dropping a disk:

- A disk is starting to fail and you must replace it
- You want to upgrade a disk
- You want to reallocate the disk to a different disk group or reallocate it to a different storage system

Note: Dropping disks from a disk group only logically deletes the disks from the disk group. It does not delete the disk contents. However, the contents are lost when you add the disk to a new disk group.

To drop one or more disks from a disk group:

1. Access the Oracle Automatic Storage Management Home page.
See ["Oracle Automatic Storage Management Home Page"](#) on page 9-1 for instructions.
2. Click the **Disk Groups** link to display the Disk Groups property page.
3. If the Oracle Automatic Storage Management Login page appears, log in as the `SYS` user, connecting as `SYSASM`. Provide the `SYS` password for the Oracle ASM instance. This password was set up when the Oracle ASM instance was created.
4. Click a link in the **Name** column to select the disk group from which you want to drop disks.
The Disk Group page appears.
5. Under the Member Disks heading, in the Select column, select the boxes for the disks that you want to drop, and then click **Remove**.
A confirmation page appears.
6. If you want to change rebalance power, or if you want set the `FORCE` option for dropping disks, do the following:
 - a. Click **Show Advanced Options**.
 - b. Optionally select the **Force** option or select a rebalance power in the Rebalance Power list.

Note: You might need the `FORCE` option if Oracle ASM cannot read from or write to the disk. See ["Dropping Disks from Disk Groups"](#) on page 4-14 for a discussion of the `FORCE` option.

7. Click **Yes** to drop the disks.

The Disk Group page returns, and displays a state of `DROPPING` for each disk being dropped.

8. Refresh the page until the dropped disks no longer appear in the disk group.

Caution: You cannot reuse or disconnect the dropped disks until the drop and rebalance operations are complete. Operations are complete when the dropped disks no longer appear in the disk group. See ["Dropping Disks from Disk Groups"](#) on page 4-14 for more information, and for other caveats for dropping disks.

Dropping Disk Groups

When you drop a disk group, you delete the disk group and optionally all of its files. You cannot drop a disk group if any of its database files are open. Oracle ASM deletes all of the disk group contents only if you specify the **including contents** option. After dropping a disk group, you can add its member disks to other disk groups or use them for other purposes.

One reason to drop a disk group is to change the redundancy level (normal, high, or external). Because you cannot change the redundancy of a disk group, you must drop the disk group and then re-create it with the proper redundancy. In this case, you must back up or move disk group data before you drop the disk group.

To drop a disk group:

1. Access the Oracle ASM Home page.

See Also: ["Oracle Automatic Storage Management Home Page"](#) on page 9-1 for instructions about how to access the Oracle Automatic Storage Management Home page

2. Click the **Disk Group** tab to display the Disk Groups page.
3. If the Oracle ASM Login page appears, log in as the `SYS` user, connecting as `SYSASM`. Provide the `SYS` password for the Oracle ASM instance. This password was set up when the Oracle ASM instance was created.
4. Click the box in the **Select** column to select a disk group that you want to drop.
5. Click **Delete** and confirmation page appears.
6. If you want to delete the disk group even if it still contains files, click **Show Advanced Options** and ensure that the **Including Contents** option is selected. If the disk group contains files and the **Including Contents** option is not selected, then you cannot drop the disk group.
7. To delete a disk group that belongs to a disk group that you cannot mount, use the `FORCE` option in the Advanced Options confirmation dialog. The `FORCE` option applies to the `DISMOUNTED` disk group only. This option enables you to delete a disk header and mark it as `FORMER`.

Caution: Use extreme care when using the `FORCE` option because the Oracle ASM instance does not verify whether the disk group is used by any other Oracle ASM instance before Oracle ASM deletes the disk group.

8. On the confirmation page, click **Yes**.

See Also: ["Dropping Disk Groups"](#) on page 4-44 for more information about dropping disks

Monitoring Disk Group Usage

Perform the following procedures to monitor disk group usage:

To view Oracle ASM disk group usage and free space:

1. Access the Oracle Automatic Storage Management Home page.

See Also: ["Oracle Automatic Storage Management Home Page"](#) on page 9-1 for instructions about how to access the home page

2. Click the **Disk Groups** link to view the Disk Groups property page.
3. If the Oracle ASM Login page appears, then log in as the `SYS` user, connecting as `SYSASM`. Provide the `SYS` password that was set up for Oracle ASM when the Oracle ASM instance was created.

The Disk Group list page, similar to the page shown in [Figure 9-3](#), displays all of the disk groups and their space usage information.

The Usable Free column displays the space in megabytes that is actually available in the disk group. The values in this column consider the redundancy level of the disk group, and exclude the space that the disk group reserves for restoring full redundancy for all files in the disk group after a disk failure.

See Also: ["Managing Capacity in Disk Groups"](#) on page 4-21 and ["Views Containing Oracle ASM Disk Group Information"](#) on page 6-1

Mounting and Dismounting Disk Groups

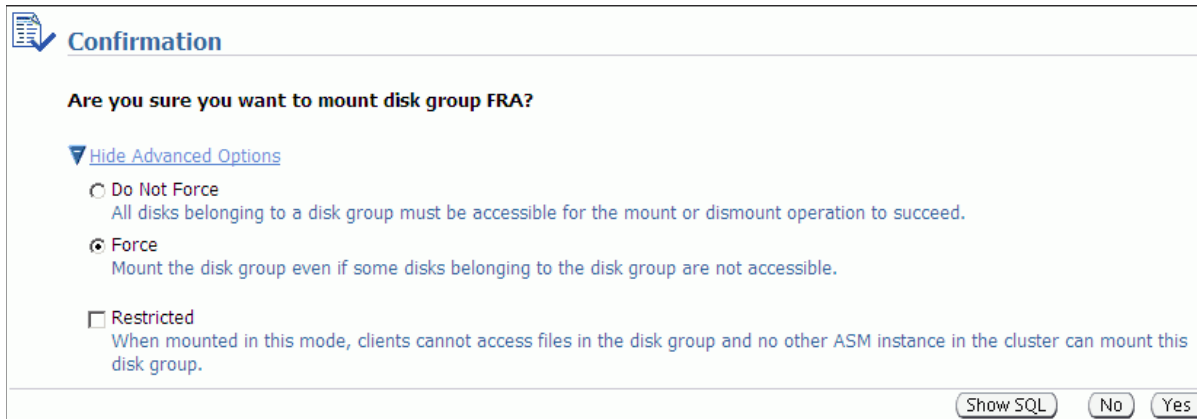
Use the `FORCE` option to mount a disk group when there are missing disks or missing failure groups. The `FORCE` option corrects configuration errors, such as incorrect values for `ASM_DISKSTRING`, without incurring unnecessary rebalance operations.

Disk groups mounted with the `FORCE` option have one or more disks offline if the disks were not available at the time of the mount. You must take corrective action to restore those devices before the time set with the `DISK_REPAIR_TIME` value expires. Failing to restore and put those disks back online within the disk repair time frame results in Oracle ASM automatically removing the disks from the disk group. This would then require a rebalance operation to restore redundancy for all of the files in the disk group. Perform the following procedures to mount a disk with the `FORCE` option:

1. From the Oracle ASM home page click the **Disk Groups** tab.
2. Select the disk group that has a disk that you must repair and click **Dismount**. Click **OK** on the Dismount confirmation dialog. The State column of the Disk Group home page should show the Status as `DISMOUNTED`.

3. Perform any needed repairs to the disk in the disk group.
4. Click the box in the Select column next to the disk that you want to force mount and click **Mount**.
5. On the Disk Mount Confirmation dialog, click the **Show Advanced Options** link.
6. On the Mount Force Confirmation dialog as shown in [Figure 9–8](#), select the **Force** option and click **Yes**. You can also click **No** to cancel the mount force operation and **Show SQL** to review the SQL that Oracle Enterprise Manager uses to perform the mount force operation.

Figure 9–8 Mount Force Confirmation Dialog



Administering Advanced Disk Group Properties

This section explains how to administer advanced disk group properties.

Perform the following steps to change the values for disk group attributes that can be altered after the disk group has been created.

1. On the Automatic Storage Disk Group List page, click a disk group name and Oracle Enterprise Manager displays the Disk Group Home page.
2. Click **Edit** under the Advanced Attributes section.

The topics contained in this section are:

- [Configuring Disk Group Compatibility Attributes](#)
- [Configuring Disk Repair Time](#)
- [Configuring Smart Scan Compatibility](#)
- [Configuring File Access Control](#)

[Figure 9–9](#) shows the advanced attributes that can be altered for an existing disk group. This screenshot only shows the top portion of the **Edit Advanced Attributes for Disk Group** Page.

Figure 9–9 Oracle Automatic Storage Management Edit Advanced Attributes for Disk Group Page

Edit Advanced Attributes for Disk Group: DATA

Disk Group Compatibility

Advancing the disk group compatibility enables the user to use new features available in the newer version. This operation can not be reversed.

Database Compatibility
The minimum software version required for a database instance to use files in this disk group (10.1 and above).

ASM Compatibility
The minimum software version required for an ASM instance to mount this disk group (10.1 and above).

ASM Volume Compatibility
The minimum software version required for an ASM Volume to use this disk group (11.2 and above).

TIP The database compatibility has to be less than or equal to the ASM compatibility. The ASM Volume compatibility can only be set when ASM compatibility is 11.2 and above.

Disk Repair Time

Specify a time period that is long enough for you to repair the disk. Keep in mind that large values, for example more than 1 day, increase the probability of data loss because the files in the disk group are protected with a lower degree of redundancy during this time period.

Disk Repair Time
Disks will be dropped after the specified time has passed.

Configuring Disk Group Compatibility Attributes

Specify a **Disk Compatibility** value for the disk group to enable your system to use the latest Oracle ASM features.

Note: You cannot reverse setting a disk compatibility value.

- **ASM Compatibility**

Specify the value that is the minimum software version required for an Oracle ASM instance to mount this disk group.

- **Database Compatibility**

Specify a value that must be less than or equal to the Oracle ASM compatibility. This is the minimum software version that is required for a database instance to use the files in the disk group. If the database compatibility value that you use is greater than the Oracle ASM compatibility, then Oracle Enterprise Manager displays an error.

Note: You cannot modify the database compatibility for an Oracle ASM cluster that is in a rolling migration state.

- **ASM Volume**

Specify a value that is 11.2 or greater to use the Oracle ADVM functionality. The ASM disk group compatibility must be set to 11.2 or greater to enable Oracle ADVM.

For information about disk group compatibility attributes, see ["Disk Group Compatibility"](#) on page 4-31.

Configuring Disk Repair Time

Specify a time period that is long enough for you to repair a disk in the disk group.

For information about disk group repair time, see ["Oracle ASM Fast Mirror Resync"](#) on page 4-27.

Configuring Smart Scan Compatibility

Specify to disable or enable smart scan compatibility for the disk group.

This attribute is only applicable to Oracle Exadata storage.

See Also: *Oracle Exadata documentation* for information about the `CELL.SMART_SCAN_CAPABLE` attribute

Configuring File Access Control

Specify to disable or enable Oracle ASM File Access Control. If Oracle ASM File Access Control is enabled, specify the default permissions for newly created files in the disk group.

For information about Oracle ASM File Access Control, see ["Managing Oracle ASM File Access Control for Disk Groups"](#) on page 4-38.

Managing Oracle ASM File Access Control with Oracle Enterprise Manager

You can manage Oracle ASM File Access Control with Oracle Enterprise Manager Integration for Oracle ASM File Access Control, shown in [Figure 9–10](#).

Figure 9–10 Oracle Automatic Storage Management Disk Group Access Control Home Page

The screenshot shows the Oracle Enterprise Manager interface for the 'Disk Group: DATA' page, specifically the 'Access Control' tab. The page is divided into two main sections: 'Users' and 'Groups'.

Users Section: Contains a description of user access, current users listed as 'grid, oracle1, oracle2', and 'Add' and 'Remove' buttons.

Groups Section: Contains a description of groups, a 'Create' button, and a table listing existing groups.

Select	Name	Owner	Members	Edit
<input type="checkbox"/>	db_users	grid	oracle1,oracle2	

To manage Oracle ASM File Access Control:

1. Access the Oracle Automatic Storage Management Home page.

See "Oracle Automatic Storage Management Home Page" on page 9-1 for instructions.

2. Click the **Disk Groups** link to display the Disk Groups property page.
3. Click a link in the **Name** column to select a disk group.
4. Click the **Access Control** tab to display the **Access Control** page.

On this page, you can add or remove users from a user group and create or delete user groups.

Figure 9–11 shows the **Create Group** page for Oracle ASM File Access Control.

Figure 9–11 Oracle Automatic Storage Management Create Group Page

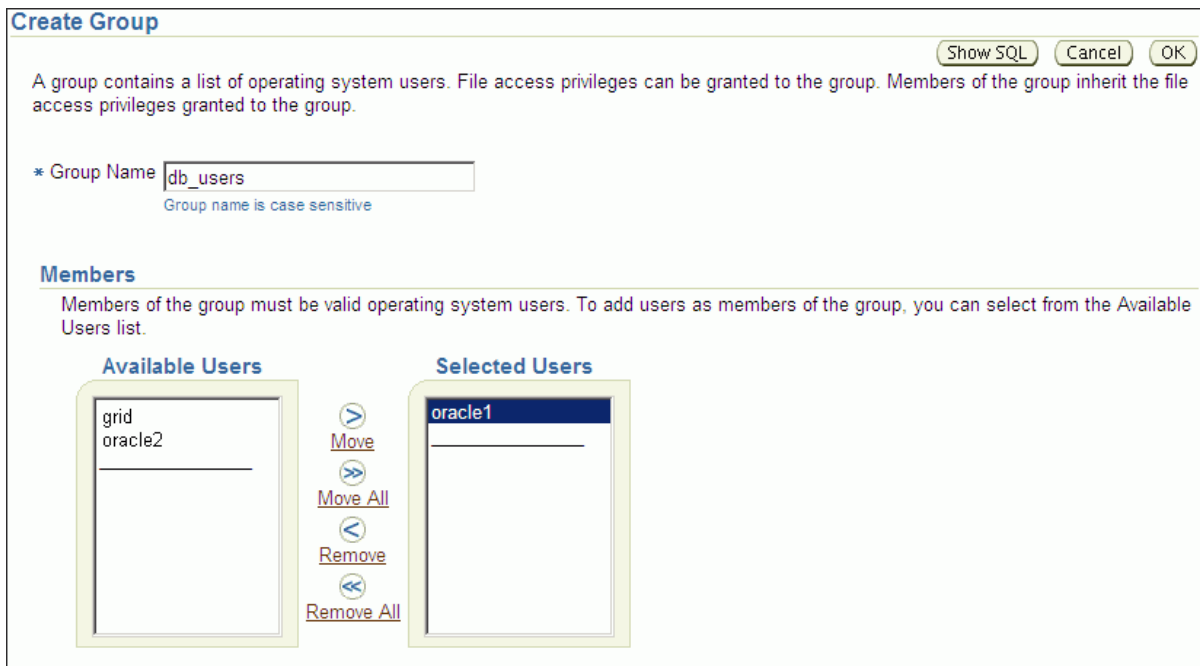
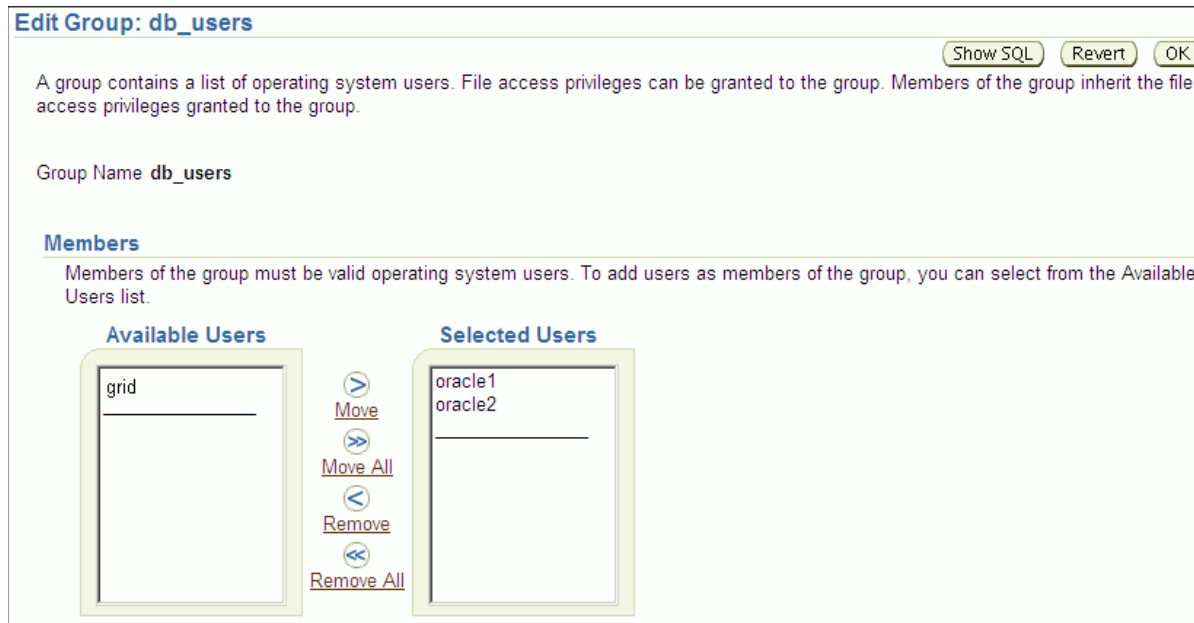


Figure 9–11 shows the **Edit Group** page for Oracle ASM File Access Control.

Figure 9–12 Oracle Automatic Storage Management Edit Group Page



For more information about Oracle ASM File Access Control, see "[Managing Oracle ASM File Access Control for Disk Groups](#)" on page 4-38.

Managing Directories, Files, and Aliases with Oracle Enterprise Manager

You can manage directories, files, and aliases with the **Files** tab of the Oracle Enterprise Manager Disk Group page, shown in [Figure 9–13](#).

Figure 9–13 Oracle Automatic Storage Management Files Tab

Disk Group: DATA

General Performance Templates **Files** Access Control Volumes

The following are the directories, files and aliases associated with the serviced databases in this disk group.

Create Alias Create Directory Rename Edit File Delete

Select All | Select None | Expand All | Collapse All

Select	Name	Physical Size (KB)	Logical Size (KB)	Primary Region	Mirror Region	Permissions			Ownership	
						Owner	Group	Other	Owner	Group
<input type="checkbox"/>	DATA									
<input type="checkbox"/>	ASM									
<input type="checkbox"/>	ASMPARAMETERFILE									
<input type="checkbox"/>	REGISTRY.253.690887395	2048	1.5	COLD	COLD	Read-write	Read-write	Read-write		
<input type="checkbox"/>	ORCL									
<input type="checkbox"/>	CONTROLFILE									
<input type="checkbox"/>	DATAFILE									
<input type="checkbox"/>	SYSAUX.257.690888309	1111040	552968	COLD	COLD	Read-write	Read-write	Read-write		
<input type="checkbox"/>	SYSTEM.256.690888307	1397760	696328	COLD	COLD	Read-write	Read-write	Read-write		
<input type="checkbox"/>	UNDOTBS1.258.690888309	107520	51208	COLD	COLD	Read-write	Read-write	Read-write		
<input type="checkbox"/>	USERS.259.690888309	12288	5128	COLD	COLD	Read-write	Read-write	Read-write		
<input type="checkbox"/>	ONLINELOG									
<input type="checkbox"/>	PARAMETERFILE									
<input type="checkbox"/>	TEMPFILE									
<input type="checkbox"/>	spfileorcl.ora	2048	3.5	COLD	COLD	Read-write	Read-write	Read-write		

To manage Oracle ASM files:

1. Access the Oracle Automatic Storage Management Home page.
See "Oracle Automatic Storage Management Home Page" on page 9-1 for instructions.
2. Click the **Disk Groups** link to display the Disk Groups property page.
3. Click a link in the **Name** column to select a disk group.
4. Click the **Files** tab to display the **Files** page.
5. Click the + to expand a database folder to view template folders.
6. Click the + to expand a templates folder to view individual files.
7. You can select a file with the **Select** box, then click **Rename** to rename the file, click **Edit** to change the properties of a file, or click **Delete** to remove a file.
You can click the name of a file to display the properties page.

REGISTRY.253.690887395: Properties OK

<p>General</p> <p>Name REGISTRY.253.690887395 Type ASMPARAMETERFILE Redundancy MIRROR Block Size (Bytes) 512 Blocks 3 Logical Size (KB) 1.5 Striped COARSE Creation Date Jun 30, 2009 9:09:55 AM PDT Modification Date Jun 30, 2009 9:00:00 AM PDT</p>	<p>Regions</p> <p>Primary COLD Mirror COLD</p> <p>Permissions</p> <p>Owner Read write Group Read write Other Read write</p> <p>Ownership</p> <p>Owner None Group None</p>
---	--

For more information about managing Oracle ASM files, directories, and aliases, see [Chapter 7, "Administering Oracle ASM Files, Directories, and Templates"](#).

Managing Disk Group Templates with Oracle Enterprise Manager

You can manage disk group templates with the Templates tab of the Disk Group page, shown in [Figure 9–14](#).

Figure 9–14 Oracle Automatic Storage Management Templates Page

Disk Group: DATA

General Performance **Templates** Files Access Control Volumes Create

Edit Delete

Select	Name	Redundancy	Striped	Internal	Primary Region	Mirror Region
<input checked="" type="radio"/>	PARAMETERFILE	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	ASMPARAMETERFILE	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	ASM_STALE	High	Coarse	Yes	COLD	COLD
<input type="radio"/>	OCRBACKUP	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	OCRFILE	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	DATAGUARDCONFIG	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	FLASHBACK	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	CHANGETRACKING	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	XTRANSPORT	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	AUTOBACKUP	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	BACKUPSET	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	TEMPFILE	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	DATAFILE	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	ONLINELOG	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	ARCHIVELOG	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	FLASHFILE	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	CONTROLFILE	High	Fine	Yes	COLD	COLD
<input type="radio"/>	DUMPSET	Mirror	Coarse	Yes	COLD	COLD
<input type="radio"/>	ASMPARAMETERBAKFILE	Mirror	Coarse	Yes	COLD	COLD

Edit Delete

To manage Oracle ASM templates:

1. Access the Oracle Automatic Storage Management Home page.
See "[Oracle Automatic Storage Management Home Page](#)" on page 9-1 for instructions.
2. Click the **Disk Groups** link to display the Disk Groups property page.
3. Click a link in the **Name** column to select a disk group.
4. Click the **Templates** tab to display the Templates page.
5. You can click **Create** to add a new template, click **Edit** to change a page, or click **Delete** to remove a template.

If you click **Create**, then the Create Template page displays and you can enter a name in the Template Name field. You can also specify a redundancy level, striped granularity, and extent regions (hot or cold).

Create Template

Show SQL Cancel OK

* Template Name

Redundancy High Mirror Unprotected

Striped Fine Coarse

Regions

Specify the regions that the primary and mirror extents should be written to based on access pattern. Place frequently accessed data in the hot region, which is in the outermost tracks. The hot region has higher bandwidth and can provide improved I/O latency

Primary Hot Cold

Mirror Hot Cold

TIP If the data is frequently accessed and mostly read only, put the primary extents in the hot region and mirror extents in the cold region. If the data is frequently accessed and updated, put both primary and mirror extents in the hot region

See "[Managing Disk Group Templates](#)" on page 7-15.

Monitoring Oracle ASM Performance with Oracle Enterprise Manager


This section describes how to perform the following performance-related Oracle ASM administrative tasks:

- [Monitoring General Oracle ASM Performance](#)
- [Checking Disk Group Integrity](#)

Monitoring General Oracle ASM Performance

You can monitor Oracle ASM performance with the Oracle Enterprise Manager Performance page, shown in [Figure 9-15](#).

Figure 9–15 Oracle Automatic Storage Management Performance Page



To monitor Oracle ASM performance:

1. Access the Oracle Automatic Storage Management Home page.
 - See Also:** "Oracle Automatic Storage Management Home Page" on page 9-1 for more information about accessing the home page
2. Click the **Performance** link to view the Performance property page.
3. Determine the view and refresh rate of the performance graphs with the **View** and **Refresh** options. Click **Refresh** to manually refresh the graphs. The refresh option has the following two intervals:
 - Manual refresh: manual, 15 second interval, 30 second interval
 - Historical data: Last 24 hours, last 7 days, last 31 days

See Also: "Performance and Scalability Considerations for Disk Groups" on page 4-30

Checking Disk Group Integrity

When you check the integrity of Oracle ASM disk groups, Oracle Enterprise Manager performs consistency checks on all of the disk group metadata. If Oracle ASM detects errors, then Oracle ASM stores information about errors in the alert log. To check disk groups, perform the following steps:

1. From the Disk Groups Administration page click **Check** and Oracle ASM displays a confirmation dialog for the check disk group operation.

Figure 9–16 Oracle Automatic Storage Management Check Diskgroup Confirmation Dialog

Check Diskgroup: DATA

The Check operation will perform consistency checks on all disk group metadata. Detected errors and information are stored in the alert log.

Check Options

Check Without Repair
Any inconsistencies will be detected and written to the alert log. But ASM will not take any automatic action to resolve them.

Check And Repair
ASM will attempt to repair any errors found during the consistency check.

2. Select **Check Without Repair** if you only want Oracle ASM to record information about disk group inconsistencies in the alert log. In this operation, Oracle ASM does not resolve the errors. Select **Check And Repair** if you want Oracle ASM to attempt to repair errors that Oracle Enterprise Manager identifies during the disk group consistency check.
3. Click **Yes** to start the disk group check operation or **No** to cancel it. You can also select **Show SQL** to review the SQL statements that the disk group check operation uses. Click **OK** to run the check operation as a command.
4. To schedule a job, click **Schedule Job**.
5. Optionally, you can modify or reenter the job name and description. Also provide host credentials and enter the scheduling options by selecting **Immediately** or **Later** and modifying the date and time as needed.
6. Click **OK** to schedule the job or **Cancel** to terminate the job creation process.

Backing Up Oracle ASM Files with Oracle Enterprise Manager

Oracle recommends that you use RMAN to back up files managed by Oracle ASM. The procedures for using RMAN are documented in *Oracle Database 2 Day DBA*.

You can also back up your Oracle ASM-managed files by performing a files related operation using the following procedures:

1. Create an alias.
2. Create a directory
3. Delete the alias and directory.
4. Rename the alias and directory

For information about migrating data with RMAN, see [Chapter 8, "Performing Oracle ASM Data Migration With RMAN"](#).

See Also: For additional information about RMAN, refer to the *Oracle Database Backup and Recovery User's Guide*.

Performing Bad Block Recovery with Oracle Enterprise Manager

The procedures that you use in Oracle Enterprise Manager to perform bad block recovery correspond to the steps that you could manually perform with the `remap`

command using the Oracle ASM Command utility (ASMCMD). Use the following procedures to perform bad block recovery:

1. Go to the Oracle Automatic Storage Management home page.
2. Click the **Disk Groups** link at the top of the page.
3. If you are not logged in, then the Log In page appears and you should log in as the SYS user, connecting as SYSASM. Provide the SYS password for the Oracle ASM instance. This password was set up when the Oracle ASM instance was created.
4. Click on a disk group name and Oracle Enterprise Manager displays the disk group home page for that disk group.
5. Click a box in the Select column to the left of a disk for which you want to recover bad blocks and click **Recover Bad Blocks**.

Note: You can only recover bad blocks for disk groups that you have configured with normal or high redundancy; you cannot use this feature for disk groups that you have configured with external redundancy.

Migrating to Oracle ASM with Oracle Enterprise Manager

This section explains how to use Oracle Enterprise Manager to migrate your database storage to use Oracle ASM.

Perform the following steps to migrate to Oracle ASM:

1. At the top of the Oracle Enterprise Manager Database Instance Home Page, click the link labeled **Server**. Oracle Enterprise Manager displays Database Server Page.
2. On the top left-hand side of the Database Server Page under the Storage heading, click the link labeled **Migrate to Oracle ASM**.
3. If prompted for Oracle ASM login credentials, then enter the user SYS, provide the SYS password that was set for the Oracle ASM instance during installation, and connect as SYSDBA. Oracle Enterprise Manager displays the Migrate Database To Oracle ASM: Migration Options Page.

For more information about authentication, see "[Authentication for Accessing Oracle ASM Instances](#)" on page 3-22.

4. On the **Migrate Database To ASM: Migration Options Page**, which is the first step in a four-step process, verify the information about the files that you want to migrate to Oracle ASM. If the target database is in ARCHIVELOG mode, then you could perform online or offline migration by selecting the corresponding radio button. Otherwise, using offline migration is the default. If you perform an online migration, then you can suspend the migration by selecting the checkbox next to **Suspend Migration Job Before Switching Database To ASM**.

Additionally, you can enter an e-mail address to which Oracle Enterprise Manager sends a message to notify you about the suspension. By default on this page, the Recovery-related Files option is selected. Deselect this option if you do not want to use it. Also verify the entry for the initialization parameter file. Lastly, the **Parallelize File Copy Operations** option is selected by default. Deselect this if you do not want Oracle Enterprise Manager to perform copy operations in parallel. Click **Next** when you complete your selections on this page and Oracle Enterprise Manager displays the **Migrate Database To ASM: Disk Group Options** page.

5. On the Disk Group Options page under Database Area, verify the location where Oracle ASM should create data files, temporary files, redo log files, and control files. Also verify the fast recovery area where Oracle ASM should create the recovery-related files such as archived redo log files, backup files, and so on. Also under fast recovery area, the option for Setup fast recovery area without Migrating Recovery-related Files is selected by default. If needed, change the selection to migrate the files. Click **Next** when you complete your selections on the Disk Group Options page.
6. If you chose to use a fast recovery area and Oracle Enterprise Manager displays a warning that the available space on the disk group is less than the specified fast recovery area size for the migration, then you can either continue by clicking **Yes**, or click **No** to stop the operation and modify your storage.
7. When you continue, Oracle Enterprise Manager displays the Schedule page on which you can specify a job name and a date and time at which you want the migration job to begin. Click **Next** when you have configured your desired settings.
8. Oracle Enterprise Manager displays the Review page on which you can review all of your selections before beginning your migration operation. To expand the tablespace details and review them, click the plus (+) sign next to Tablespaces and Files To Be Migrated. Click **Back** to modify any of your selections or **Submit Job** to begin the migration. Click **Cancel** to cancel the operation.

Caution: If your migration job fails, then do not start your database until you complete a recovery operation. Otherwise, you may corrupt your database.

9. After you submit your job, Oracle Enterprise Manager should display a confirmation dialog indicating that the migration job was successful. Click **View Status** on the confirmation dialog to examine the progress of your migration job.

After the migration job has completed, the Oracle ASM links on the Oracle Enterprise Manager page may not display immediately. You can wait for the links to display after a subsequent refresh, or restart the Oracle Enterprise Manager agent as follows:

```
$ emctl stop dbconsole
$ emctl start dbconsole
```

Oracle ASM Support Workbench

Oracle Enterprise Manager provides Oracle ASM Support Workbench to monitor Oracle ASM alerts and incidents.

To access Support Workbench for Oracle ASM:

1. Click the **Software and Support** tab on the database home page.
2. Click **Support Workbench** under the **Support** section on the **Software and Support** page.
3. Click **Support Workbench**(*ASM_instance_name*) under the **Related Links** section on the **Support Workbench** page.

Figure 9–17 shows the Oracle ASM Support Workbench page with the **Problems** tab displayed. You can view information about current and past problems in this tab. On this page, you have options to view additional details and create a package that you can send to Oracle Support Services.

Figure 9–17 Oracle Automatic Storage Management Support Workbench Problems

Support Workbench

Page Refreshed July 22, 2009 9:07:47 AM PDT

Problems (1) [Checker Findings \(0\)](#) [Packages \(0\)](#)

New Problems in Last 24 Hours 1 All Active Problems 1 All Problems 1
 New Incidents in Last 24 Hours 2 All Active Incidents 1 All Incidents 2

View Search [Advanced Search](#)

[Select All](#) | [Select None](#) | [Show All Details](#) | [Hide All Details](#)

Select	Details	ID	Description	Number Of Incidents	Last Incident	Last Comment	Active	Packaged	SR#
<input type="checkbox"/>	Hide	1	ORA 7445 [__kernel_vsyscall()+2]	2	July 22, 2009 9:07:23 AM PDT		Yes	No	

Incidents (2)

922	ORA-7445 [__kernel_vsyscall()+2] [SIGSEGV] [ADDR:0x1B61] [PC:0x193402] [unk ...	July 22, 2009 9:07:23 AM PDT
921	ORA-7445 [__kernel_vsyscall()+2] [SIGSEGV] [ADDR:0x1B61] [PC:0x572402] [unk ...	July 22, 2009 9:04:40 AM PDT

To create a package, select an incident and click **Package** on the Support Workbench Problems page. Support Workbench guides you through the packaging process.

Figure 9–18 shows step one the Oracle ASM Support Workbench **Create New Package** wizard.

Figure 9–18 Oracle Automatic Storage Management Support Workbench Create Package

Quick Packaging: Create New Package

Step 1 of 4

Target +ASM_dadvmn0652.us.oracle.com Logged in As sys

Problems Selected ORA 7445 [__kernel_vsyscall()+2]

Use quick packaging to generate an upload file for a single problem and send it to Oracle with default options. If Oracle Configuration Manager is set up, the upload file will still be created but it will not be sent to Oracle.

* Package Name

Package Description

Send to Oracle Support Yes No

My Oracle Support Username

My Oracle Support Password

Customer Support Identifier (CSI)

Country

Create new Service Request (SR) Yes No

Figure 9–19 shows the Oracle ASM Support Workbench page with the **Packages** tab displayed. On this tab, you can view and edit information about incident packages created for problems.

Figure 9–19 Oracle Automatic Storage Management Support Workbench Check Packages

Support Workbench Page Refreshed July 22, 2009 9:24:03 AM PDT [Refresh](#)

[Problems \(1\)](#) [Checker Findings \(0\)](#) [Packages \(1\)](#)

Active 0 Contents Finalized 0
 Upload File 1 Sent to Oracle 0
 Generated

Search [Go](#)

[Delete](#)

[Select All](#) | [Select None](#)

Select	Name	Status	Description	Main Problem Keys	Created
<input type="checkbox"/>	ORA7445_20090722090952	Upload File Generated	Oracle ASM test problem	ORA 7445 [__kernel_vsycall()+2]	July 22, 2009 9:21:13 AM PDT

You can view additional details for a problem by selecting a problem on the **Problems** tab, and then clicking on **View**. The problem details page displays as shown in Figure 9–20.

Figure 9–20 Oracle Automatic Storage Management Support Workbench Problem Details

Problem Details: ORA 7445 [__kernel_vsycall()+2] Page Refreshed July 22, 2009 9:25:47 AM PDT [Refresh](#)

Summary

SR# .. [Edit](#)
 Bug# .. [Edit](#)
 Active **Yes**
 Packaged **Yes**
 Number of Incidents 2
 First Incident [July 22, 2009 9:04:40 AM PDT](#)

Last Dumped Incident

Timestamp [July 22, 2009 9:07:23 AM PDT](#)
 Incident Source **System Generated**
 Impact
 Checkers Run 0
 Checker Findings 0

Investigate and Resolve

[Go to My Oracle Support](#) [Quick Package](#)

[Self Service](#) [Oracle Support](#)

Diagnose

[Alert Log](#)
[Related Problems Across Topology](#)
[Diagnostics for Last Dumped Incident](#)
[Go to My Oracle Support and Research](#)

[Incidents](#) [Activity Log](#)

Status [Open Incidents](#) Data Dumped [Yes](#) [Go](#)

[View](#) [Close](#)

[Select All](#) | [Select None](#) | [Show All Details](#) | [Hide All Details](#)

Select	Details	ID	Description	Data Dumped	Active	Status	Timestamp
<input type="checkbox"/>	Show	922	ORA-7445 [__kernel_vsycall()+2] [SIGSEGV] [ADDR:0x1B61] [PC:0x193402] [unknown code]	Yes	Yes	Ready	July 22, 2009 9:07:23 AM PDT
<input type="checkbox"/>	Show	921	ORA-7445 [__kernel_vsycall()+2] [SIGSEGV] [ADDR:0x1B61] [PC:0x572402] [unknown code]	Yes	Yes	Ready	July 22, 2009 9:04:40 AM PDT

On the Problem Details page, you can access additional diagnostic tools. You can also create a quick package to send to Oracle Support Services.

See Also:

- *Oracle Database Administrator's Guide* for information about viewing problems with the Oracle Enterprise Manager Support Workbench
- *Oracle Database 2 Day DBA* for information about accessing the Oracle Enterprise Manager Support Workbench home page and viewing problems using the Oracle Enterprise Manager Support Workbench

Oracle ASM: Oracle by Example Series

Oracle by Example (OBE) has a series on the *Oracle Database 2 Day DBA* book. This OBE steps you through Oracle ASM tasks and includes annotated screen shots.

To view the Oracle ASM OBE, point your browser to the following location:

http://www.oracle.com/technology/obe/10gr2_2day_dba/asm/asm.htm

Administering Oracle ACFS With Oracle Enterprise Manager

This chapter describes how to administer Oracle Automatic Storage Management Cluster File System (Oracle ACFS) using Oracle Enterprise Manager.

This chapter contains the following topics:

- [Managing Oracle ACFS with Oracle Enterprise Manager](#)
- [Managing Oracle ACFS Snapshots with Oracle Enterprise Manager](#)

All Oracle ASM administration tasks begin with the Oracle Automatic Storage Management Home page in Oracle Enterprise Manager. For information about the Oracle ASM home page and how to access that page, see ["Oracle Automatic Storage Management Home Page"](#) on page 9-1.

For information about using command-line tools to administer Oracle ACFS file systems, see [Chapter 13, "Oracle ACFS Command-Line Tools"](#). For information about the basic steps for creating Oracle ACFS file systems, see ["Basic Steps to Manage an Oracle ACFS"](#) on page 13-1.

Managing Oracle ACFS with Oracle Enterprise Manager

This discusses how to manage Oracle ACFS volumes and file systems with Oracle Enterprise Manager.

This section contains the following topics:

- [Creating Oracle ACFS Volumes and File Systems](#)
- [Viewing and Modifying Oracle ACFS Volumes and File Systems](#)

Creating Oracle ACFS Volumes and File Systems

To create an Oracle ACFS volume, perform the following steps.

1. Access the **Disk Groups** page from the Oracle ASM home page.
2. Click the **Volumes** link in the **General** tab of the **Disk Group** page. See [Figure 10–6](#) on page 10-5.
3. Click **Create** in the **Volumes** tab of the **Disk Group** page.

Enter the volume name (maximum 11 characters), disk group name that contains the volume, and the initial size of the volume. You can also select the redundancy setting for the volume and the region settings for primary and mirror extents.

[Figure 10–1](#) shows the **Oracle ASM Create ASM Volume** page.

Figure 10–1 Oracle Automatic Storage Management Create ASM Volume Page

Create ASM Volume Show Command Cancel Ok

Disk Group FRA Disk Group Usable Free Space (GB) 3.2343
 Volume Allocation Unit (MB) 256

* Volume

* Size

Volume size will be adjusted to be in multiples of the Volume Allocation Unit.

Redundancy

Stripe Width (KB) 128
 Stripe Columns 4

Regions

Specify the regions that the primary and mirror extents should be written to based on access pattern. Place frequently accessed data in the hot region, which is in the outermost tracks. The hot region has higher bandwidth and can provide improved I/O latency

Primary Hot Cold
 Mirror Hot Cold

TIP If the data is frequently accessed and mostly read only, put the primary extents in the hot region and mirror extents in the cold region. If the data is frequently accessed and updated, put both primary and mirror extents in the hot region

To create an Oracle ACFS file system on a volume in a disk group, perform the following steps.

1. Access the **Volumes** tab of the **Disk Group** page. See [Figure 10–6](#) on page 10-5.
2. Select an existing volume in the disk group and click **Create ASM Cluster File System**.

[Figure 10–2](#) shows the ASM Create Oracle ACFS file system page.

The compatibility parameters `COMPATIBLE.ASM` and `COMPATIBLE.ADVM` must be set to 11.2 or higher for the disk group. See ["Disk Group Compatibility Attributes"](#) on page 4-33.

3. On the ASM Create Oracle ACFS file system page, enter the information to create a file system. You can optionally register and mount the file system.

You need to enter the volume name. You can enter an optional volume label.

You can optionally choose to register and mount the file system when it is created. Select a mount point from available directories.

Enter the host credentials needed to execute the command when prompted. To register or mount a file system, you need root or administrator privileges. For information about the registry process, see ["About the Oracle ACFS Mount Registry"](#) on page 5-6.

After a directory has been selected, click **Show Command** to have Oracle Enterprise Manager generate the commands that you can execute at an operating system prompt. Examples of commands on a Linux system are:

```
/sbin/mkfs -t acfs -b 4k /dev/asm/volume1-361
/sbin/acfsutil registry -f -a /dev/asm/volume1-361 /oracle/acfsmounts/acfs2
/bin/mount -t acfs -rw /dev/asm/volume1-361 /oracle/acfsmounts/acfs2
```

You can also generate the commands to register or mount an Oracle ACFS file system on the **ASM Cluster File System** tab. See [Figure 10–3](#) on page 10-3.

Figure 10–2 Oracle Automatic Storage Management Create Oracle ACFS Page

Create ASM Cluster File System Show Command Cancel Ok

Creating ASM Cluster File System creates the on-disk structure. Specify a volume device that is not currently used by any file system to prevent loss of data on the volume. The file system needs to be mounted before it is available for use.

* Volume Device Create ASM Volume

Volume Label

Register ASM Cluster File System Mount Point

Registering the ASM Cluster File System mount point makes the mount point available for future mount operations.

Mount Point Help

The entire directory path must already exist on the host file system.

Related Links

[Open Telnet Session](#)

Viewing and Modifying Oracle ACFS Volumes and File Systems

To view information about Oracle ACFS file systems, click the **ASM Cluster File System** link in the ASM home page.

Figure 10–3 shows the **ASM Cluster File System** tab in Oracle Enterprise Manager. This tab lists all of the Oracle ACFS associated with the Oracle ASM instance.

On this page, you can choose to mount, dismount, delete, create snapshot, view content, register, and deregister a selected file system. In addition, you can create a file system, mount all file systems, or dismount all file systems.

For each Oracle ACFS, the columns provide information for Mount Point, Availability, State, Snapshots, Used (%), Used (GB), Size (GB), Allocated Space, Redundancy, Volume, and Disk Group. Redundancy, Volume, and Disk Group columns are on the far right of the page and are not shown in illustration. The mount point, snapshots, volume, and disk group are provided as links to further information.

Figure 10–3 Oracle Automatic Storage Management Oracle ACFS Page

Home Performance Disk Groups Configuration Users ASM Cluster File System									
Mount Dismount Delete Create Snapshot View Content Register Deregister Actions Resize									
Select	Mount Point	Availability	State	Snapshots	Used (%)	Used (GB)	Size (GB)	Allocated Space	Redundancy
<input checked="" type="radio"/>	/u01/app/acfsmounts/acfs1		Mounted	0	<div style="width: 4.89%;"></div> 4.89	0.0979	2		
<input type="radio"/>	/u01/app/acfsmounts/acfs2		Mounted	0	<div style="width: 3.73%;"></div> 3.73	0.0374	1		

TIP Allocated Space is determined by the size and redundancy of the File System

To view information about a specific volume, click the volume name in the **Volume** column on the **ASM Cluster File System** page.

Figure 10–4 shows the **General** tab of the **ASM Volumes** page.

Figure 10–4 Oracle Automatic Storage Management General Tab of the Volumes Page

ASM Volume: VOLUME1

General Edit Ok

Disk Group DATA Disk Group Usable Free Space (GB) 1.833

State **ENABLED**
 Usage ACFS
 Volume Device /dev/asm/volume1-413
 Mount Point [/u01/app/acfsmounts/acfs1](#)
 Size (GB) 2
 Allocated Space (GB) 4.0029
 Redundancy **Mirror**
 Stripe Width (KB) 128
 Stripe Columns 4
 Volume Allocation Unit (MB) 256
Volume Allocation Unit is the unit by which volume grows or shrinks
 Primary Region Cold
 Mirror Region Cold
 Creation Time Jun 30, 2009 10:03:52 AM PDT
 Last Modification Time Jun 30, 2009 10:00:00 AM PDT

To view information about a file system, click the link in the **Mount Point** column on the **ASM Cluster File System** page.

Figure 10–5 shows the **General** tab of the ASM Oracle ACFS page.

Figure 10–5 Oracle Automatic Storage Management Oracle ACFS General Page

General Snapshots Data Retrieved July 2, 2009 11:22:59 AM PDT Refresh

General

Current Status **Up**
 Up Since Jun 30, 2009 10:14:06 AM PDT
 Check and Repair **No**
 Volume [VOLUME1](#)
 Volume Device /dev/asm/volume1-413

Current Space Usage

Size (GB)	2.00
Used (GB)	0.10
Free (GB)	1.90
Used (%)	4.89
Allocated Space (GB)	4.00
Space Used by Snapshots (MB)	0.00

Daily Space Usage History

Size (GB) vs. Time (June 2009)

■ Total Size

Serviced Hosts

Host Name	ACFS Status	Available Since	State	ASM Instance
dadvmn0652.us.oracle.com ↑		Jun 30, 2009 10:14:06 AM PDT	Mounted	+ASM_dadvmn0652.us.oracle.com ↑

Alerts

Severity	Category	Name	Message	Alert Triggered
(No alerts)				

To view information about a disk group for an Oracle ACFS, click the disk group name in the **Disk Group** column on the **ASM Cluster File System** page. The **General** tab of the **Disk Group** page displays as shown in Figure 9–4 on page 9-5.

To view information about volumes in a disk group, click the **Volumes** tab at the top of the **Disk Group** page.

Figure 10–6 shows the Oracle ACFS Volume page.

Figure 10–6 Oracle Automatic Storage Management Oracle ACFS Volume Page

Select	Volume	Volume Device	State	Usage	Mount Point	Size (GB)	Allocated Space (GB)	Redundancy
<input type="checkbox"/>	VOLUME1	/dev/asm/volume1-413	ENABLED	ACFS	/u01/app/acfsmounts/acfs1	2	4.0029	Mirror

Managing Oracle ACFS Snapshots with Oracle Enterprise Manager

To view information about existing snapshots or create snapshots for a file system, perform the following steps.

1. Click on the **ASM Cluster File System** tab on the Oracle ASM home page
2. Click on a mount point link in the **Mount Point** column as shown in Figure 10–3 on page 10-3
3. Click on the **Snapshots** tab, shown in Figure 10–7.

Optionally, you can click a number link for existing snapshots in the **Snapshots** column on the **ASM Cluster File System** page to display the **Snapshots** page.

On the **Snapshots** page, you can create snapshots or search for and display specific snapshots. To search for a snapshot, enter a name with optional wildcard characters in the search field then click **Search**.

Figure 10–7 Oracle Automatic Storage Management Oracle ACFS Snapshots Tab

ASM Cluster File System: /u01/app/acfsmounts/acfs1

General | **Snapshots**

Summary
 A snapshot is a space efficient point-in-time copy of the file system. It can be used as a backup source to recover accidentally deleted or modified files, or for data mining and report applications.

Number of Snapshots 4
 Last Snapshot Creation Time Jul 2, 2009 2:24:21 PM PDT
 Space Used by Snapshots (MB) 0.1289
 File System Free Space (GB) 1.9021

Snapshots

Search
 To narrow down the search results or to search for a specific snapshot, enter a search string in the Snapshot name field and click Search. To display all the ACFS snapshots for the selected ACFS filesystem, clear the search field and click Search.

Name

TIP To return exact or case-sensitive matches, double quote the search strings. You can use the wildcard symbol (% , *) in a double-quoted string.

[Select All](#) | [Select None](#)

Select	Name	Creation Time
<input type="checkbox"/>	snapshot_20090702_140518	Jul 2, 2009 2:05:57 PM PDT
<input type="checkbox"/>	snapshot_20090702_141102	Jul 2, 2009 2:11:05 PM PDT
<input type="checkbox"/>	snapshot_20090702_141449	Jul 2, 2009 2:14:52 PM PDT
<input type="checkbox"/>	snapshot_20090702_142110	Jul 2, 2009 2:24:21 PM PDT

To create a snapshot, perform the following steps.

1. Click **Create** in the **Snapshots** page.

Figure 10–8 shows the **Create Snapshot** page.

2. Complete the information on the page.

Accept the default snapshot name, or provide a name. Optionally, you can choose to delete the oldest snapshot.

3. When you have completed the screen, you can click **OK** to run the command, or click **Show Command** to view the generated command.

For example, the following is a generated command for creating a snapshot:

```
/sbin/acfsutil snap create "snapshot_20090702_142135" /oracle/acfsmounts/acfs1
```

To run the generated command, you need the appropriate privileges on the host computer. Oracle Enterprise Manager prompts you to enter host credentials if they have not been set up.

You can also open a Telnet session from the **Create Snapshot** page if you want to manually enter the generated operating system command.

Figure 10–8 Oracle Automatic Storage Management Create Snapshot Page

Create Snapshot on ASM Cluster File System: /u01/app/acfsmounts/acfs1 Show Command Cancel Ok

Number of Snapshots 3 Last Snapshot Name [snapshot_20090702_141449](#)
 Total Free Space (GB) 1.9021 Last Snapshot Creation Time Jul 2, 2009 2:14:52 PM PDT

The Create Snapshot operation creates a point-in-time copy of the ASM Cluster File System. Snapshots are always mounted and available. The snapshots are created in the .ACFS/snaps directory of the file system with the snapshot name as the directory name.

* Snapshot Name
 Delete the oldest snapshot (61 remaining before the maximum limit is reached)
 Oldest Snapshot Name [snapshot_20090702_140518](#)
 Oldest Snapshot Creation Time Jul 2, 2009 2:05:57 PM PDT

Related Links
[Open Telnet Session](#)

To drill down in a snapshot directory, click on the snapshot name in the **Snapshots** page to display the **Search and List Directory** page.

Figure 10–9 shows the **Search and List Directory** page for a specific snapshot directory. The directory contains the snapshots of the file system

Figure 10–9 Oracle Automatic Storage Management Search and List Directory Page

Search and List: Directory Close

Host dadvmn0652
 Host Username **grid** Change

Current Directory [/ > u01 > app > acfsmounts > acfs1 > .ACFS > snaps > snapshot_20090702_142110 >](#)

Search

Search in Directory
 Directory path only, no wildcard.

File or Directory Name
 This search is case sensitive. Use * or % as wildcard.

Go

Name	Type	Owner	Group	Size(KB)	Last Modified
.	Directory	grid	asmadmin	4	Jul 2, 2009 2:24:21 PM
..	Directory	root	root	4	Jul 2, 2009 2:24:21 PM
.ACFS	Directory	root	root	4	Jun 30, 2009 10:14:06 AM
lost+found	File	root	root	0	Dec 31, 1969 4:00:00 PM
mytestfile	File	grid	asmadmin	0	Jul 2, 2009 2:10:01 PM

Oracle ASM Configuration Assistant

Oracle ASM Configuration Assistant (ASMCA) supports installing and configuring ASM instances, disk groups, volumes, and Oracle Automatic Storage Management Cluster File System (Oracle ACFS). In addition, you can use the ASMCA command-line interface as a non-GUI utility.

This chapter discusses the following topics:

- [Oracle ASM Configuration Assistant GUI Tool](#)
- [Oracle ASM Configuration Assistant Command-Line Interface](#)

Oracle ASM Configuration Assistant GUI Tool

This section contains the following topics:

- [Starting Oracle ASM Configuration Assistant](#)
- [Managing Oracle ASM Instances With Oracle ASM Configuration Assistant](#)
- [Upgrading an Oracle ASM Instance with Oracle ASM Configuration Assistant](#)
- [Managing Disk Groups With Oracle ASM Configuration Assistant](#)
- [Managing Oracle ADVM Volumes With Oracle ASM Configuration Assistant](#)
- [Managing Oracle ACFS File Systems With Oracle ASM Configuration Assistant](#)
- [Creating an Oracle ACFS File System for a Database Home](#)

Starting Oracle ASM Configuration Assistant

To start ASMCA, enter `asmca` at a command prompt inside the Oracle grid infrastructure home. For example:

```
$ asmca
```

The GUI tool displays and attempts to connect to Oracle ASM instance identified with the Oracle system identifier (SID) set to `+ASM`.

- If ASMCA detects an Oracle ASM instance that is not running, ASMCA prompts whether you want to start the instance.
- If ASMCA detects an Oracle ASM instance from a previous release, ASMCA prompts whether you want to upgrade the instance. Oracle recommends that you respond `No` this prompt.

An Oracle ASM instance should be upgraded with Oracle Universal Installer (OUI). OUI automatically defaults to upgrade mode when it detects an Oracle

ASM instance at a previous release level. See ["Upgrading an Oracle ASM Instance With Oracle Universal Installer"](#) on page 3-17.

If you wish to continue the upgrade process with ASMCA, see ["Upgrading an Oracle ASM Instance with Oracle ASM Configuration Assistant"](#) on page 11-3. If you need help with upgrading an Oracle ASM instance, please contact Oracle Support Services.

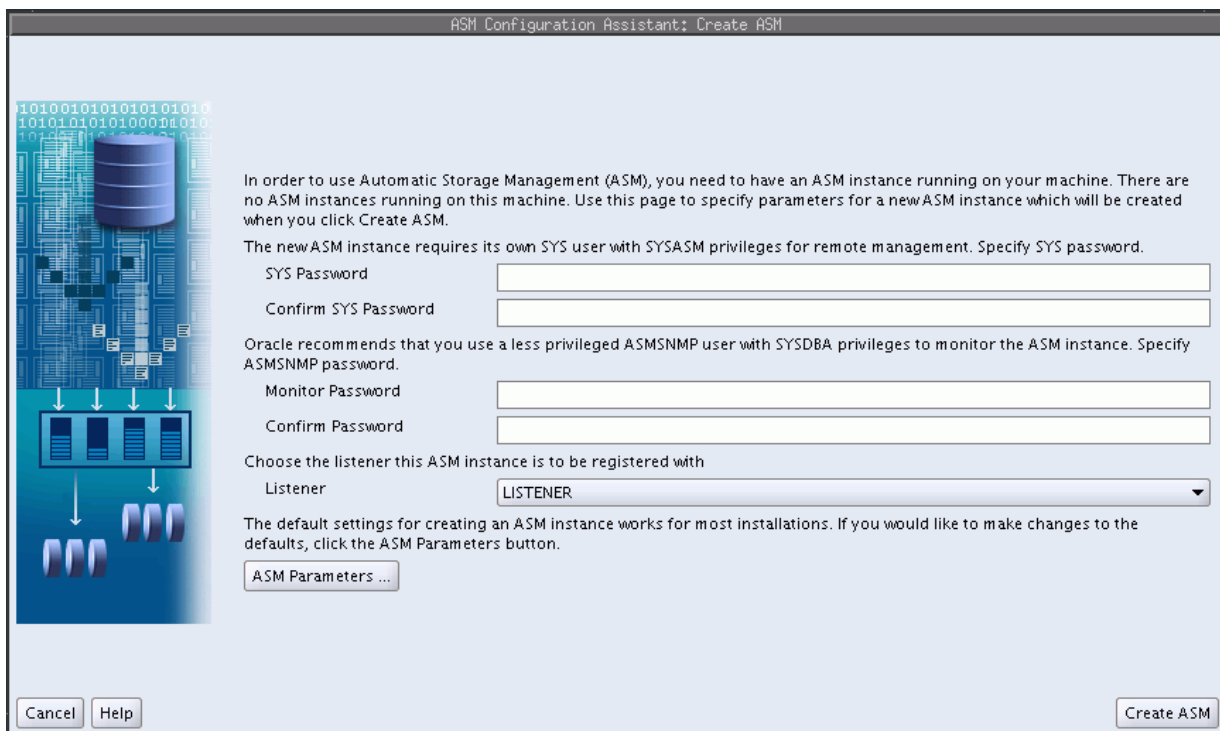
To start the command-line ASMCA, see ["Oracle ASM Configuration Assistant Command-Line Interface"](#) on page 11-14.

Managing Oracle ASM Instances With Oracle ASM Configuration Assistant

Oracle ASM Configuration Assistant enables you to create a new Oracle ASM instance if an Oracle ASM instance is not found on the host, as shown in [Figure 11-1](#).

On the **Create ASM** page, you enter the SYS password for the administrator of the Oracle ASM instance and the password for the user that can monitor this Oracle ASM instance. You can select the listener with which you want to register this Oracle ASM instance. Click **ASM Parameters** to update the initialization parameters for the Oracle ASM instance.

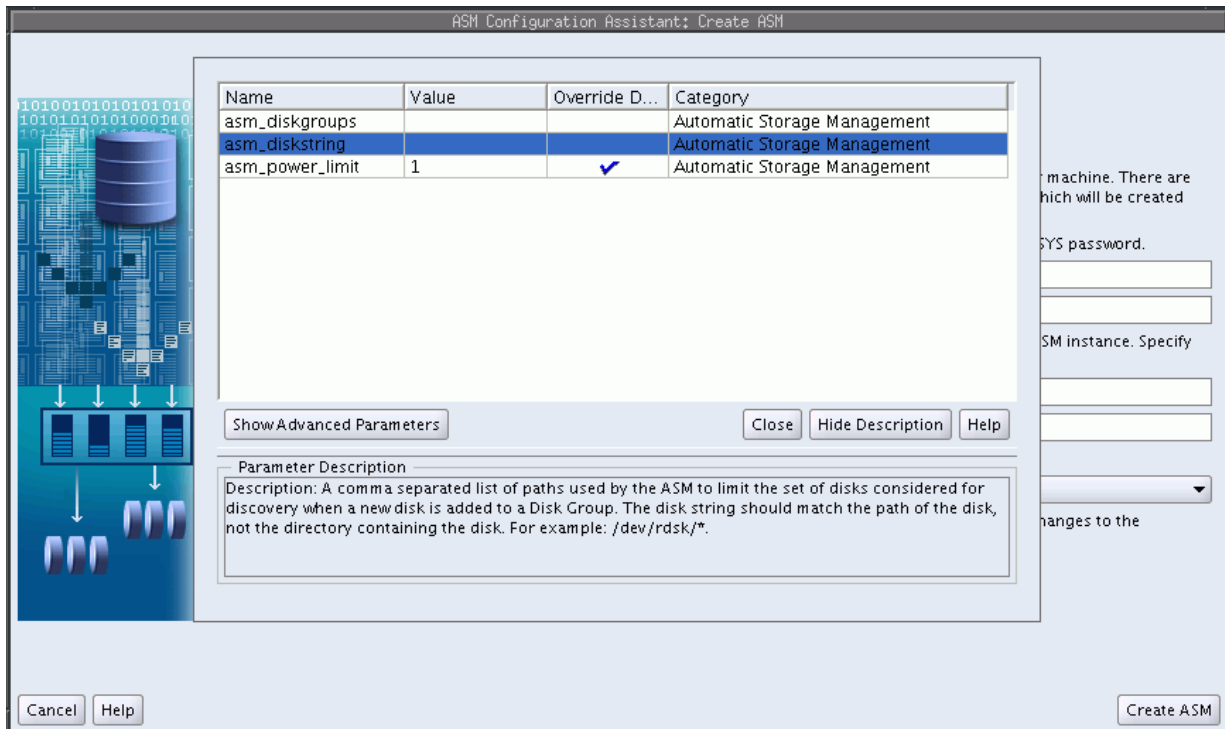
Figure 11-1 Oracle ASM Configuration Assistant Create ASM Page



When you can click **ASM Parameters** in the **Create ASM** page, you can update initialization parameters for an Oracle ASM instance on the Oracle ASM Configuration Assistant page, as shown in [Figure 11-2](#).

For information about Oracle ASM instance initialization parameters, see ["Configuring Initialization Parameters for an Oracle ASM Instance"](#) on page 3-2.

Figure 11–2 Oracle ASM Configuration Assistant Create ASM Page



Upgrading an Oracle ASM Instance with Oracle ASM Configuration Assistant

Oracle ASM Configuration Assistant also enables you to upgrade an existing Oracle ASM instance to the current software level. However, the recommended practice is to upgrade an Oracle ASM instance with Oracle Universal Installer (OUI). OUI automatically defaults to upgrade mode when it detects an Oracle ASM instance at a previous release level. See ["Upgrading an Oracle ASM Instance With Oracle Universal Installer"](#) on page 3-17.

The following procedure describes how to upgrade an Oracle ASM instance from 11g Release 1 (11.1) to 11g Release 2 (11.2). In this scenario:

- The Oracle ASM and Oracle Database instances 11g Release 1 (11.1) exist in separate homes.
 - The Oracle grid infrastructure software 11g Release 2 (11.2) is installed in a separate home.
 - The Oracle ASM instance 11g Release 2 (11.2) is set up as an Oracle Restart (single-instance) configuration.
1. Shut down the Oracle Enterprise Manager agent, Oracle Database instance, and the listener the in the older database home.
 - Run `emctl stop dbconsole` to stop the Oracle Enterprise Manager agent.
 - Connect to the database instance with SQL*Plus as a privileged user and issue the `SHUTDOWN` command.
 - Run `lsnrctl` and enter the `STOP` command to stop the listener.

See Also:

- Oracle Enterprise Manager manuals and online help for information about starting and stopping the Oracle Enterprise Manager agent
- *Oracle Database Administrator's Guide* for more information about starting up and shutting down Oracle instances
- *Oracle Database Net Services Administrator's Guide* for information about configuring a listener

2. Run the Oracle Universal Installer (OUI) and select the **Install Oracle grid infrastructure Software Only** option.
3. Complete the screens in the OUI installer and run the scripts as prompted by the OUI installer.

For example, on Linux you must run the `root.sh` script as the `root` user.

```
# GRID_HOME/grid/perl/bin/root.sh
```

When you run the `root.sh` script on Linux, note the instructions in the output of that script. You must run different scripts as `root` to configure either a standalone (single-instance) or cluster configuration after OUI has completed.

See Also: *Oracle Grid Infrastructure Installation Guide* for information about installing Oracle grid infrastructure

4. After the OUI has completed, perform the procedures that are described in the output of the `root.sh` script that you ran during the installation.

For example, for a standalone (single-instance) configuration on Linux, run the following script as `root` to upgrade the necessary services, including Oracle Cluster Synchronization Services (CSS).

```
# GRID_HOME/grid/perl/bin/perl -I/GRID_HOME/perl/lib -I/GRID_HOME/crs/install  
GRID_HOME/crs/install/roothas.pl
```

For an upgrade, you must update the node list by running OUI as the installation user in the Oracle grid infrastructure home.

```
$ GRID_HOME/oui/bin/runInstaller -updateNodeList -silent -local CRS=false  
ORACLE_HOME=ASM_11G_R1_HOME
```

```
$ GRID_HOME/oui/bin/runInstaller -updateNodeList -silent -local CRS=true  
ORACLE_HOME=GRID_HOME
```

5. Ensure that the current Oracle Cluster Synchronization Services (CSS) are running in the home where the Oracle ASM instance is located. For example:

On Linux run the `crsctl check css` command. If the CSS daemon is not running, run `localconfig add` as the `root` user.

On Windows, you can use the `crsctl` and `localconfig` commands or the Windows Services GUI tool. To use Windows Services GUI tool, double-click the Services icon in the Windows Control Panel and locate the `OracleCSService` service. The status of `OracleCSService` should be `Started` and its startup type should be `Automatic`.

6. Run the Network Configuration Assistant (NETCA) to configure the listener.

Start Network Configuration Assistant (NETCA) in the Oracle grid infrastructure home with `netca`. Follow the prompts in the wizard to add and configure the listener in the Oracle grid infrastructure home.

7. Run ASMCA from the Oracle grid infrastructure home to complete the upgrade.

For information about starting ASMCA, see "[Starting Oracle ASM Configuration Assistant](#)" on page 11-1. If the Oracle ASM instance has been shut down during the Oracle grid infrastructure configuration, start that instance when ASMCA prompts you to start that instance.

When ASMCA detects an earlier version of Oracle ASM, the assistant automatically determines the location of that instance and the location of the Oracle grid infrastructure home, as shown in [Figure 11-3](#).

Enter a password for the `ASMSNMP` user, then select **Upgrade ASM** to complete the upgrade. When the process has completed, you can exit from ASMCA.

When upgrading an Oracle ASM instance to the current software level, ASMCA moves an Oracle ASM instance prior to 11g Release 2 (11.2) to the Oracle grid infrastructure home and updates the appropriate environmental variables.

8. Ensure that the Oracle Database instance and Oracle Enterprise Manager agent are running in the database home 11g Release 1 (11.1).

Start the following services if necessary.

- Connect to the Oracle Database instance with SQL*Plus as a privileged user and issue the `STARTUP` command.
- Restart the Oracle Enterprise Manager agent with `emctl start dbconsole`.

Note: The procedure described in this section upgrades the Oracle ASM instance only. Oracle Database, and Oracle Enterprise Manager, will not have the latest features. To upgrade Oracle Database, see *Oracle Database Upgrade Guide*.

Figure 11–3 Oracle ASM Configuration Assistant Upgrade ASM Page

After upgrading an Oracle ASM instance, you can copy or move an Oracle ASM. After copying or moving the SPFILE, you must restart the instance with the SPFILE in the new location if you want to use that SPFILE. For information, see ["Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File"](#) on page 3-4.

See Also:

- *Oracle Database Upgrade Guide* for information upgrading an Oracle ASM instance
- *Oracle Grid Infrastructure Installation Guide* for information about installing Oracle grid infrastructure and performing a rolling upgrade of Oracle ASM

Managing Disk Groups With Oracle ASM Configuration Assistant

Oracle ASM Configuration Assistant enables you to configure or create Oracle ASM disk groups with the **Configure ASM Disk Groups** tab, as shown in [Figure 11–4](#). The disk group tab displays selected details about the disk group, such as name, size, free space, usable space, redundancy level, and state of the disk group.


See Also: *Oracle Database Reference* for descriptions of disk group information displayed in the V\$ASM_DISKGROUP view

You can right click on a selected disk group in the screen to display the configuration options menu. The menu options enable you to:

- Add disks to the disk group
- Edit the disk group attributes
- Manage templates for the disk group

- Create an Oracle ACFS file system on the disk group for a database home
- Dismount and mount the disk group
- Drop the disk group

Figure 11–4 Oracle ASM Configuration Assistant Configure Disk Groups Page



There are additional options for creating a disk group, mounting all disk groups, and dismounting all disk groups.

Oracle ASM Configuration Assistant enables you to create an Oracle ASM disk group by clicking **Create** on the Oracle ASM Configuration Assistant disk group page. The page for creating a disk group is shown in [Figure 11–5](#).

Figure 11–5 Oracle ASM Configuration Assistant Create Disk Groups Page

Disk Group Name

Redundancy

Redundancy is achieved by storing multiple copies of the data on different failure groups. Normal redundancy needs disks from at least two different failure groups, and high redundancy from at least three different failure groups.

High Normal External (None)

Select Member Disks

Show Eligible Show All

Quorum failure groups are used to store voting files in extended clusters and do not contain any user data. It requires ASM compatibility of 11.2 or higher.

<input type="checkbox"/>	Disk Path	Header Status	Disk Name	Size (MB)	Failure Group	Quorum
<input type="checkbox"/>	/devices/diskd1	CANDIDATE		1024		<input type="checkbox"/>
<input type="checkbox"/>	/devices/diskd2	CANDIDATE		1024		<input type="checkbox"/>
<input type="checkbox"/>	/devices/diskd3	CANDIDATE		1024		<input type="checkbox"/>
<input type="checkbox"/>	/devices/diskj1	CANDIDATE		1024		<input type="checkbox"/>
<input type="checkbox"/>	/devices/diskj2	CANDIDATE		1024		<input type="checkbox"/>
<input type="checkbox"/>	/devices/diskj3	CANDIDATE		1024		<input type="checkbox"/>

Note: If you do not see the disks which you believe are available, check Disk Discovery Path and read/write permissions on the disks. The Disk Discovery Path limits set of disks considered for discovery.

Disk Discovery Path: /devices/disk*

Click on the Show Advanced Options button to change the diskgroup attributes. Diskgroup compatibility attributes may need to be modified based on the usage of diskgroup for different versions of databases or ASM Cluster File Systems.

On this page, you can enter a name for the disk group, select the redundancy level, change the disk discovery path to locate disks in a different directory, identify the disks you want to be in the disk group, enter a name for a failure group, and specify if the failure group is a quorum type. A quorum failure group is a special type of failure group and disks in these failure groups do not contain user data and are not considered when determining redundancy requirements. For information about quorum failure groups, see ["Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups"](#) on page 4-10.

The default disk group compatibility settings are 11.2 for Oracle ASM compatibility, 10.1 for database compatibility, and no value for Oracle ADVM compatibility. For information about disk group compatibility attributes, see ["Disk Group Compatibility"](#) on page 4-31. For information about disk discovery, see ["Oracle ASM Disk Discovery"](#) on page 4-18 and ["ASM_DISKSTRING"](#) on page 3-7.

You can choose to show advanced options when creating a disk group, as shown in [Figure 11–6](#). The advanced options page enables you to set additional options, such as disk group compatibility settings and allocation unit (AU) size. When you are finished with the create disk group specifications, click **OK**.

Note: The disk group compatibility settings can only be advanced; you cannot revert to a lower compatibility setting. After you advance the disk group compatibility settings, you cannot change the values back to the previous settings. See ["Overview of Disk Group Compatibility"](#) on page 4-32

Figure 11–6 Oracle ASM Configuration Assistant Create Disk Groups Page - Advanced Options

Disk Group Name

Redundancy
 Redundancy is achieved by storing multiple copies of the data on different failure groups. Normal redundancy needs disks from at least two different failure groups, and high redundancy from at least three different failure groups.

High Normal External (None)

Select Member Disks
 Show Eligible Show All

Quorum failure groups are used to store voting files in extended clusters and do not contain any user data. It requires ASM compatibility of 11.2 or higher.

<input checked="" type="checkbox"/>	Disk Path	Header Status	Disk Name	Size (MB)	Failure Group	Quorum
<input checked="" type="checkbox"/>	/devices/diskd1	CANDIDATE		1024		<input type="checkbox"/>
<input checked="" type="checkbox"/>	/devices/diskd2	CANDIDATE		1024		<input type="checkbox"/>
<input checked="" type="checkbox"/>	/devices/diskd3	CANDIDATE		1024		<input type="checkbox"/>
<input checked="" type="checkbox"/>	/devices/diskj1	CANDIDATE		1024		<input type="checkbox"/>
<input checked="" type="checkbox"/>	/devices/diskj2	CANDIDATE		1024		<input type="checkbox"/>
<input checked="" type="checkbox"/>	/devices/diskj3	CANDIDATE		1024		<input type="checkbox"/>

Note: If you do not see the disks which you believe are available, check Disk Discovery Path and read/write permissions on the disks. The Disk Discovery Path limits set of disks considered for discovery.

Disk Discovery Path: /devices/disk*

Diskgroup Attributes
 An allocation unit (AU) is the fundamental unit in which contiguous disk space is allocated to ASM files. ASM file extent size is a multiple of AUs. The AU size cannot be modified later.

Allocation Unit Size (MB)

Specify minimum software versions for ASM, Database and ASM volumes that this diskgroup need to be compatible with.

ASM Compatibility

Database Compatibility

ADVM Compatibility

Refer Oracle Database Storage Administrator Guide for more details on Compatibility matrix.

For more information about managing disk groups, see [Chapter 4, "Administering Oracle ASM Disk Groups"](#).


Managing Oracle ADVM Volumes With Oracle ASM Configuration Assistant

Oracle ASM Configuration Assistant enables you to create or configure an Oracle ADVM volume, as shown in [Figure 11–7](#).

Right click a selected volume in the Oracle ASM Configuration Assistant screen to display the configuration options menu. The menu options include:

- Enable a disabled volume
- Disable an enabled volume
- Resize the volume
- Delete the volume

There are options for creating a volume, enabling all volumes, and disabling all volumes.

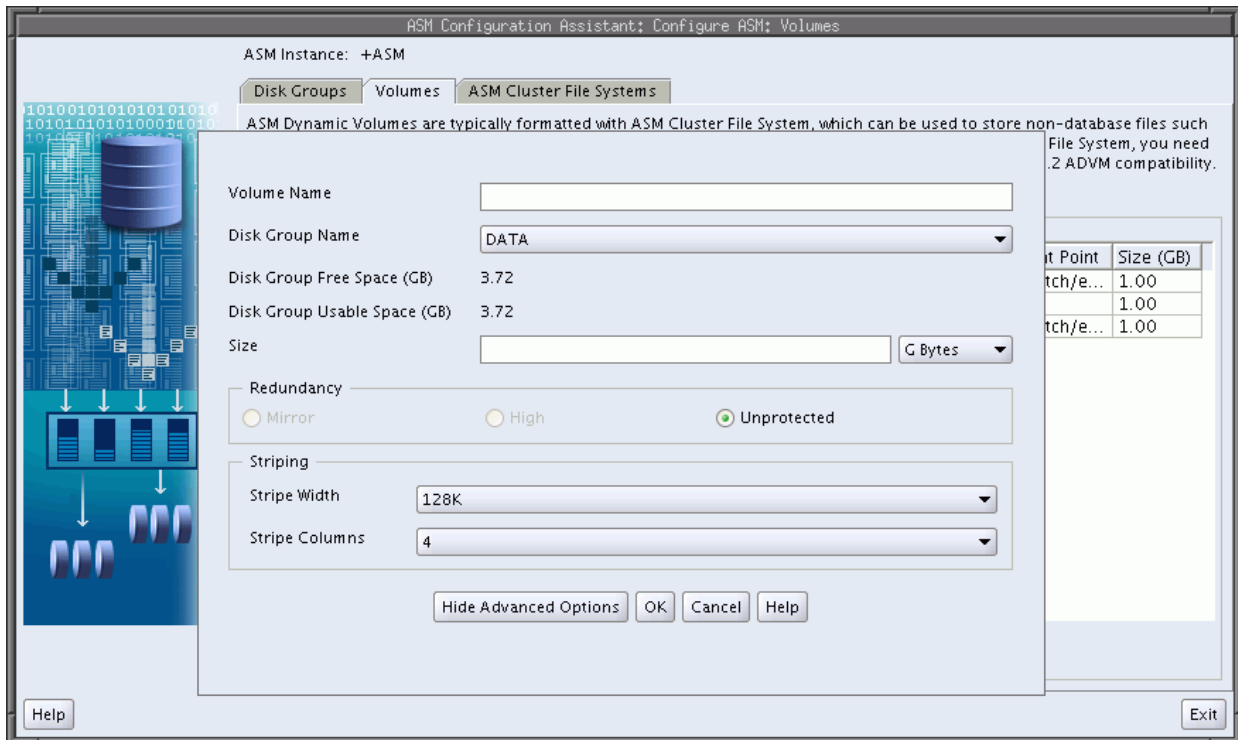
Figure 11–7 Oracle ASM Configuration Assistant Configure ASM Volumes Page

Oracle ASM Configuration Assistant enables you to create Oracle ADVM volumes, as shown in [Figure 11–7](#).

Click **Create** to create a new volume. You must provide a unique volume name for the existing Oracle ASM disk group that you select from the list. The volume name can be a maximum of 11 alphanumeric characters; dashes are not allowed. The first character must be alphabetic.

Specify the size for the volume. You can optionally choose to display advanced options for creating a volume which enables you to specify the redundancy level and striping characteristics.

The compatibility parameters `COMPATIBLE.ASM` and `COMPATIBLE.ADVM` must be set to 11.2 or higher for the disk group. See "[Disk Group Compatibility Attributes](#)" on page 4-33.

Figure 11–8 Oracle ASM Configuration Assistant Create an ASM Volume Page

For information about Oracle ACFS, see [Chapter 5, "Introduction to Oracle ACFS"](#).

Managing Oracle ACFS File Systems With Oracle ASM Configuration Assistant

Oracle ASM Configuration Assistant enables you to create or configure an Oracle ACFS file system, as shown in [Figure 11–9](#).


To configure an existing file system, right click a selected file system in the Oracle ASM Configuration Assistant screen to display the configuration options menu.

The menu options include:

- Show mount command for the file system
- Show dismount command for the file system
- Register mount point for the file system
- Deregister mount point for the file system
- Resize the file system
- Delete the file system

Some of the commands require root privileges, such as mounting a file system. Oracle ASM Configuration Assistant generates the command for you to run manually as root or as a privileged user.

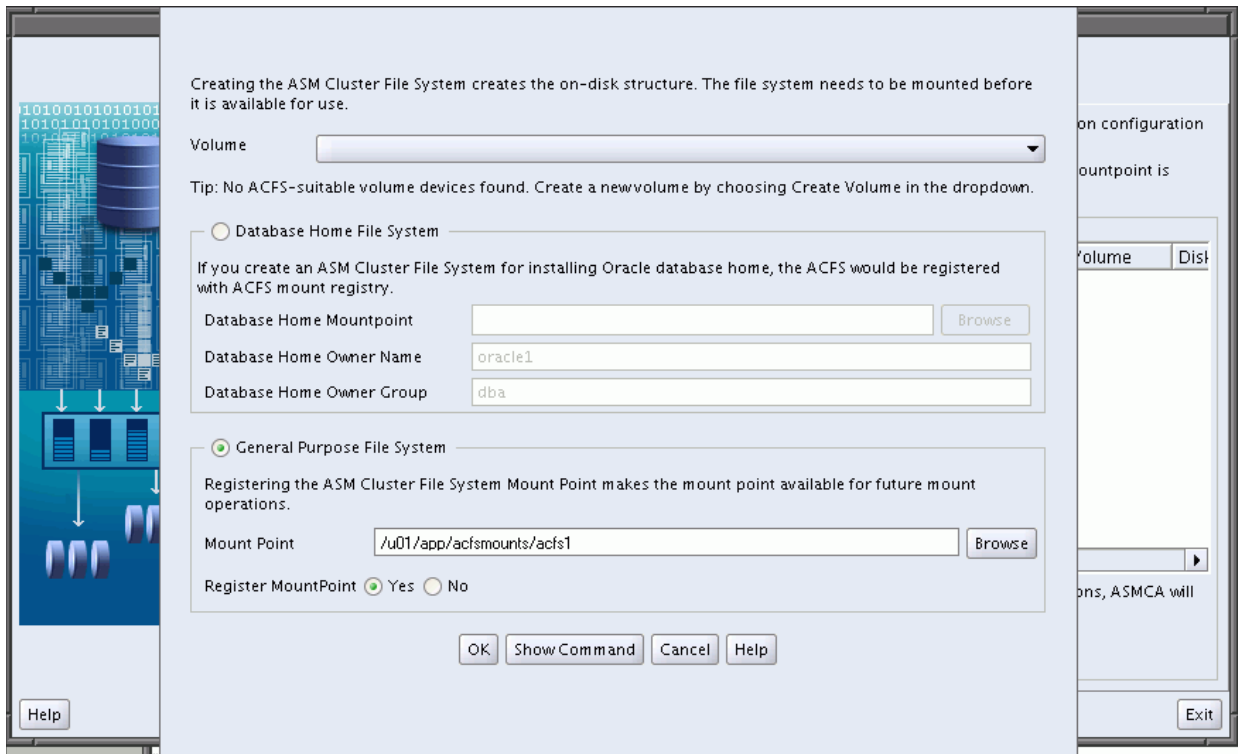
There are buttons for Create, Show Mount All, and Show Dismount All commands.

Figure 11–9 Oracle ASM Configuration Assistant Configure Oracle ACFS Page

Oracle ASM Configuration Assistant enables you to create an Oracle ACFS file system, as shown in [Figure 11–9](#).

Select **Create**, then select an existing Oracle ADVM volume from the list and determine the mount point, show in [Figure 11–10](#). The mount point must be an existing directory. The file system must be mounted to make it available. The mount command must be manually executed as root or the Windows Administrator at an operating system prompt. For information about mounting an Oracle ACFS file system, see "[mount](#)" on page 13-10 for Linux or "[acfsmountvol](#)" on page 13-17 for Windows.

You can also choose to create an Oracle ACFS file system for an Oracle Database. For more information about creating a file system for a database home, see "[Creating an Oracle ACFS File System for a Database Home](#)" on page 11-13.

Figure 11–10 Oracle ASM Configuration Assistant Create Oracle ACFS Page

For information about Oracle ACFS, see [Chapter 5, "Introduction to Oracle ACFS"](#). For a summary of the basic steps for creating an Oracle ACFS file system, see ["Basic Steps to Manage an Oracle ACFS"](#) on page 13-1.

Creating an Oracle ACFS File System for a Database Home

If you want to create an Oracle ACFS file system for a database home, you can select the **Create ACFS for Database Home** from the disk group configuration options menu, as shown in [Figure 11–4](#) on page 11-7. The **Create ACFS Hosted Database Home** dialog displays as shown in [Figure 11–11](#).

In this dialog, you need to enter:

- Database Home Volume Name

This is the name of the Oracle ADVM volume you want to create. For information about Oracle ADVM volumes, see ["About Oracle ACFS"](#) on page 5-4.
- Database Home Mount Point

This the mount point for the file system where you want to install the database home. The file system that contains the database home should not be located under the Oracle grid infrastructure base (ORACLE_BASE for grid) directory. For information about mount points and database homes, see ["About Oracle ACFS and Oracle Database Homes"](#) on page 5-4.
- Database Home Size in gigabytes (GB)

The default is 6 GB and the minimum size allowed.
- Database Home Owner Name

This is the operating system name of the user that installs the database and owns the software in the database home.

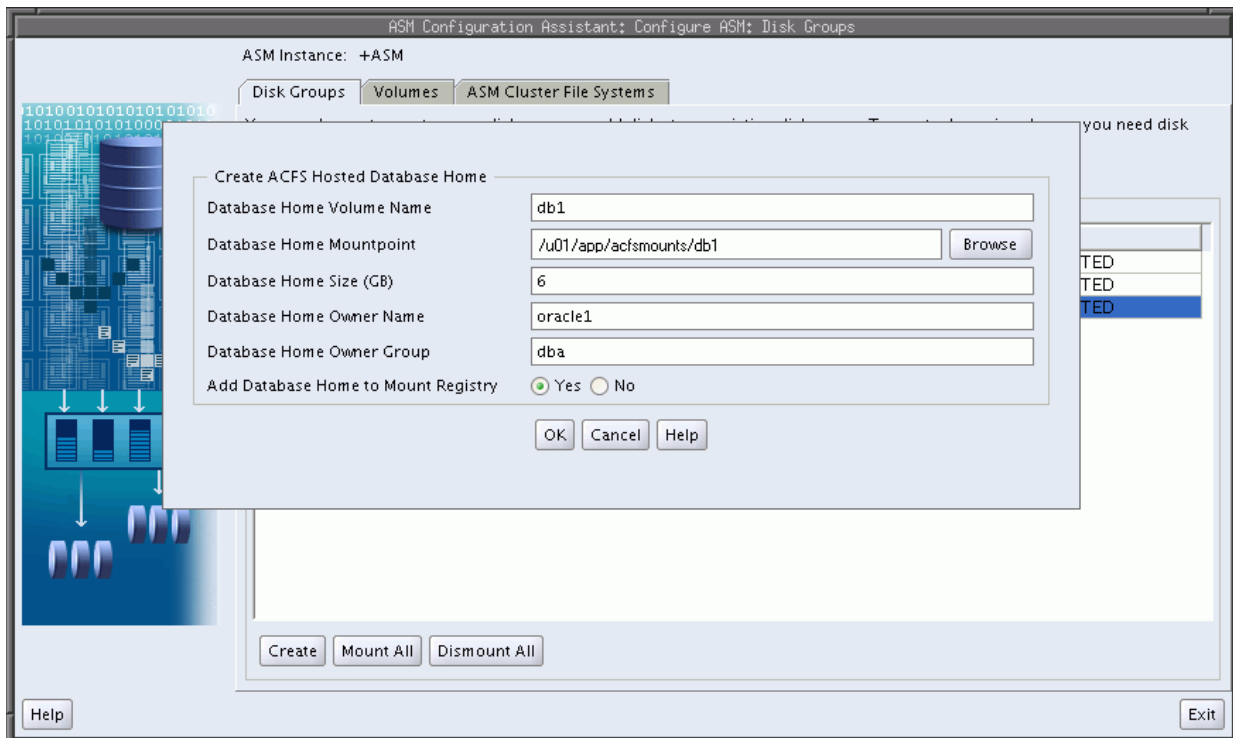
- Database Home Owner Group

This is the operating system group of the owner of the database home.

The mount point must be an existing directory. The file system must be mounted to make it available. The mount command must be manually executed as `root` or the Windows Administrator at an operating system prompt. For information about mounting an Oracle ACFS file system, see ["mount"](#) on page 13-10 for Linux or ["acfsmountvol"](#) on page 13-17 for Windows.

For information on the mount registry, see ["About the Oracle ACFS Mount Registry"](#) on page 5-6.

Figure 11–11 Oracle ASM Configuration Assistant Create Oracle ACFS Hosted Database Home Page



Oracle ASM Configuration Assistant Command-Line Interface

The ASMCA command-line interface provides a non-GUI support for configuring Oracle ASM disk groups, volumes, and Oracle ACFS.

- [Running ASMCA Command-Line](#)
- [ASMCA Commands](#)

Running ASMCA Command-Line

The syntax for running the ASMCA command-line tool is:

```
asmca -silent command_name [options]
```

[Table 11–1](#) contains the options available with the `asmca` command-line tool.

Table 11-1 Options for the *asmca* command-line tool

Option	Description
<code>-silent</code>	Specifies to run <i>asmca</i> in command-line mode.
<code>command_name</code>	<p>Specifies command that you want to run. The command can be any of the following:</p> <ul style="list-style-type: none"> -configureASM -upgradeASM -configureParameter -deleteASM -createDiskGroup -addDisk -createVolume -createACFS <p>See "ASMCA Commands" on page 11-15.</p>
<code>options</code>	<p>Specifies the options to run for a command. The options vary by command. See the individual commands described in "ASMCA Commands" on page 11-15 for the available options.</p> <p>The following list describes the options common for multiple commands.</p> <ul style="list-style-type: none"> ■ <code>-help</code> Displays help for running ASMCA in silent mode. ■ <code>-asmsnmpPassword <i>asmsnmp_password</i></code> The Oracle ASM monitor user (ASMSNMP) password. This is optional and only required when creating and migrating Oracle ASM instance. ■ <code>-sysAsmPassword <i>sysasm_password</i></code> The SYSASM password for Oracle ASM instance. This option is required.

The following is an example of the `-help` option.

```
$ asmca -help
asmca [-silent] [-sysAsmPassword <SYS user password>] [-asmsnmpPassword <ASMSNMP password>]
{<command> <options>}
Please refer to the manual for details.
You can enter one of the following commands:
...
```

ASMCA Commands

This section describes the commands that can be run with ASMCA command-line.

- [Configure an Oracle ASM Instance](#)
- [Upgrade an Oracle ASM Instance](#)
- [Configure Parameters for an Oracle ASM Instance](#)
- [Delete an Oracle ASM Instance](#)
- [Create a Disk Group](#)
- [Add a Disk to a Disk Group](#)
- [Create a Volume](#)
- [Create an Oracle ACFS File System](#)

Configure an Oracle ASM Instance

`-configureASM` configures or creates an Oracle ASM instance.

Syntax

```
asmca -silent
      -configureASM
        [-sysAsmPassword sysasm_password]
        [-asmsnmpPassword asm_monitor_password]
        [{ -param paramname=paramvalue, ... } ]
        [{ -diskString discovery_path
          -diskGroupName diskgroup
          -diskList disk_list
          [-redundancy { HIGH | NORMAL | EXTERNAL } ] ] ]
```

Table 11–2 contains the options available with the `-configureASM` command.

Table 11–2 Options for the `-configureASM` command

Option	Description
<code>-param <i>paramname=paramvalue</i>, ...</code>	Specifies a list of initialization parameters with values for the Oracle ASM instance.
<code>-diskString <i>discovery_path</i></code>	Specifies the Oracle ASM discovery string for locating disks.
<code>-diskGroupName <i>diskgroup</i></code>	Specifies the name of the Oracle ASM disk group that you want to create.
<code>-diskList <i>disk_list</i></code>	Specifies a comma separated list of disk names to be added to the disk group.
<code>-redundancy { HIGH NORMAL EXTERNAL }</code>	Specifies the redundancy setting for the disk group.

For additional options that are common to multiple commands, see Table 11–1 on page 11-15.

Examples

To configure an Oracle ASM instance:

Example 11–1 Using `asmca -silent -configureASM`

```
asmca -silent
      -configureASM
        -diskString= '/devices/disk*'
        -sysAsmPassword my_sysasm_password
        -asmsnmpPassword my_asmsnmp_passwd
```

Upgrade an Oracle ASM Instance

`-upgradeASM` upgrades an Oracle ASM instance from a previous release to the current software release.

Upgrading an Oracle ASM instance with ASMCA only updates the Oracle ASM instance software and does not configure the Oracle Clusterware or listener.

An Oracle ASM instance should be upgraded with Oracle Universal Installer (OUI). OUI automatically defaults to upgrade mode when it detects an Oracle ASM instance

at a previous release level. See ["Upgrading an Oracle ASM Instance With Oracle Universal Installer"](#) on page 3-17.

Syntax

```
asmca -silent
      -upgradeASMInstance
      [ -asmsnmpPassword asmsnmp_password ]
```

The options available with the `-upgradeASMInstance` command are common to multiple commands. For a description of those options, see [Table 11-1](#) on page 11-15.

Examples

To upgrade an Oracle ASM instance:

Example 11-2 Using `asmca -silent -upgradeASM`

```
asmca -silent
      -upgradeASM
      -asmsnmpPassword my_asmsnmp_passwd
```

Configure Parameters for an Oracle ASM Instance

`-configureParameter` modifies the parameters of an Oracle ASM instance.

Syntax

```
asmca -silent
      -configureParameter
      [ { -param paramname=paramvalue, ... } ]
      [ -sysAsmPassword sysasm_password ]
```

[Table 11-3](#) contains the options available with the `-configureParameter` command.

Table 11-3 Options for the `-configureParameter` command

Option	Description
<code>-param <i>paramname=paramvalue</i>, ...</code>	Specifies a list of initialization parameters with values for the Oracle ASM instance.

For additional options that are common to multiple commands, see [Table 11-1](#) on page 11-15.

Examples

To configure parameters for an Oracle ASM instance:

Example 11-3 Using `asmca -silent -configureParameter`

```
asmca -silent
      -configureParameter
      (-param
        asm_power_limit=3,
        asm_disk_string='/devices/disk*')
      -sysAsmPassword my_sysasm_password
```

Delete an Oracle ASM Instance

`-deleteASM` deletes an Oracle ASM instance. The operation stops and an error is raised if there is a database using the Oracle ASM instance.

Syntax

```
asmca -silent
      -deleteASM
        [-sysAsmPassword sysasm_password]
        [-force ]
        [-dropDiskGroups ]
```

Table 11-4 contains the options available with the `-deleteASM` command.

Table 11-4 Options for the `-deleteASM` command

Option	Description
<code>-force</code>	Forces the deletion of an Oracle ASM instance.
<code>-dropDiskGroups</code>	Drops the disk groups on this instance.

For additional options that are common to multiple commands, see Table 11-1 on page 11-15.

Examples

To delete an Oracle ASM instance:

Example 11-4 Using `asmca -silent -deleteASM`

```
asmca -silent
      -deleteASM
        -sysAsmPassword my_sysasm_password
        -force
```

Create a Disk Group

`-createDiskGroup` creates disk groups in an Oracle ASM instance.

Syntax

```
asmca -silent
      -createDiskGroup
        [-diskString discovery_path]
        { -diskGroupName diskgroup
          { -disk disk_path [-diskName disk] |
            -diskList disk_list
              [-diskSize disk_size_MB]
              [-failuregroup failure_group]
              [-force|-noforce]
              [-quorum|-noquorum] }
          [-redundancy { HIGH | NORMAL | EXTERNAL } ] }
        [-au_size value]
        [-compatible.asm value]
        [-compatible.rdbms value]
        [-compatible.advm value]
        [-sysAsmPassword sysasm_password]
```

Table 11-5 contains the options available with the `-createDiskGroup` command.

Table 11-5 Options for the `-createDiskGroup` command

Option	Description
<code>-diskString discovery_path</code>	Specifies the Oracle ASM discovery string for locating disks.
<code>-diskGroupName diskgroup</code>	Specifies the name of the Oracle ASM disk group that you want to create.
<code>-disk disk_path</code>	Specifies the full path of the disk device.
<code>-diskName disk</code>	Specifies a user-supplied disk name.
<code>-diskList disk_list</code>	Specifies a comma separated list of disk names to be added to the disk group.
<code>-diskSize disk_size_MB</code>	Specifies the size of the disk to use for Oracle ASM storage.
<code>-failuregroup failure_group</code>	Specifies the failure group name.
<code>-force -noforce</code>	Specifies where to force the use of the disk.
<code>-quorum -noquorum</code>	Specifies a quorum or standard failure group.
<code>-redundancy { HIGH NORMAL EXTERNAL }</code>	Specifies the redundancy setting for the Oracle ASM disk group.
<code>-au_size</code>	Specifies the allocation unit size (1,2,4,8,16,32,64)
<code>-compatible.asm</code>	Specifies the attribute setting (10.1 to 11.2)
<code>-compatible.rdbms</code>	Specifies the attribute setting (10.1 to 11.2)
<code>-compatible.advm</code>	Specifies the attribute setting (empty or 11.2)

For additional options that are common to multiple commands, see Table 11-1 on page 11-15.

The default disk group compatibility settings are 11.2 for Oracle ASM compatibility, 10.1 for database compatibility, and no value for Oracle ADVM compatibility. For information about disk group compatibility attributes, see "Disk Group Compatibility" on page 4-31.

Examples

To create disk groups in an Oracle ASM instance:

Example 11-5 Using `asmca -silent -createDiskGroup`

```
$ asmca -silent -createDiskGroup
  -diskGroupName mynewdg
    -disk '/devices/diske*'
    -disk '/devices/diskk*'
  -redundancy NORMAL
  -au_size 64
  -compatible.asm '11.2.0.0.0'
  -compatible.rdbms '11.2.0.0.0'
  -compatible.advm '11.2.0.0.0'
```

DiskGroup mynewdg created successfully.

Add a Disk to a Disk Group

`-addDisk` add disks to an existing disk group in an Oracle ASM instance.

Syntax

```
asmca -silent
      -addDisk
        [ -diskString discovery_path ]
          { -diskGroupName diskgroup
            { -disk disk_path [-diskName disk] |
              -diskList disk_list
                [-diskSize disk_size_MB]
                [-failuregroup failure_group]
                [-force | -noforce]
                [-quorum | -noquorum] }
            [-sysAsmPassword sysasm_password ]
```

Table 11–6 contains the options available with the `-addDisk` command.

Table 11–6 Options for the `-addDisk` command

Option	Description
<code>-diskString <i>discovery_path</i></code>	Specifies the Oracle ASM discovery string for locating disks.
<code>-diskGroupName <i>diskgroup</i></code>	Specifies the name of the Oracle ASM disk group that you want to update.
<code>-diskList <i>disk_list</i></code>	Specifies a comma separated list of disk names to be added to the disk group.
<code>-diskSize <i>disk_size_MB</i></code>	Specifies the size of the disk to use for Oracle ASM storage.
<code>-failuregroup <i>failure_group</i></code>	Specifies the failure group name.
<code>-force -noforce</code>	Specifies where to force the use of the disk.
<code>-quorum -noquorum</code>	Specifies a quorum or standard failure group.

For additional options that are common to multiple commands, see Table 11–1 on page 11-15.

Examples

To add disks to an existing disk group:

Example 11–6 Using `asmca -silent -addDisk`

```
asmca -silent
      -addDisk
        -diskGroupName mynewdg
          -disk '/devices/disk11'
          -disk '/devices/disk12'
```

Disks added successfully to diskgroup mynewdg.

Create a Volume

`-createVolume` creates Oracle ADVM volumes.

Syntax

```
asmca -silent
  -createVolume
    { -volumeName volume_name
      -volumeDiskGroupName diskgroup
      -volumeSizeGB size_GB
      [ -volumeRedundancy { INHERIT | HIGH | NORMAL | EXTERNAL } ] ... }
  [-sysAsmPassword sysasm_password]
```

Table 11–7 contains the options available with the `-createVolume` command.

Table 11–7 Options for the `-createVolume` command

Option	Description
<code>-volumeName <i>volume_name</i></code>	Specifies the name of the volume that you want to create.
<code>-volumeDiskGroupName <i>diskgroup</i></code>	Specifies the name of the disk group where you want to create the volume.
<code>-volumeSizeGB <i>size_GB</i></code>	Specifies the size of the volume in Gigabytes..
<code>-volumeRedundancy { INHERIT HIGH NORMAL EXTERNAL }</code>	Specifies the redundancy setting for the volume..

For additional options that are common to multiple commands, see Table 11–1 on page 11-15.

Examples

To create an Oracle ADVM volume:

Example 11–7 Using `asmca -silent -createVolume`

```
$ asmca -silent
  -createVolume
    -volumeName volume1
    -volumeDiskGroup mynewdg
    -volumeSizeGB 1
```

Volume *volume1* created successfully.

Create an Oracle ACFS File System

`-createACFS` creates an Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

This command does not mount the Oracle ACFS file system. For information about mounting an Oracle ACFS file system, see "[mount](#)" on page 13-10 for Linux environments and "[acfsmountvol](#)" on page 13-17 for Windows environments.

Syntax

```
asmca -silent
  -createACFS
    { -acfsVolumeDevice volume_device }
  [-sysAsmPassword sysasm_password]
```

Table 11–8 contains the options available with the `-createACFS` command.

Table 11–8 Options for the `-createACFS` command

Option	Description
<code>-acfsVolumeDevice volume_device</code>	Specifies the name of the Oracle ADVM volume device.

For additional options that are common to multiple commands, see [Table 11–1](#) on page 11-15.

Examples

To create an Oracle ACFS file system:

Example 11–8 Using `asmca -silent -createACFS`

```
$ asmca -silent
  -createACFS
    -acfsVolumeDevice /dev/asm/volume1-457
```

ACFS creation on /dev/asm/volume1-457 completed successfully.
The ACFS should be mounted for use.

Oracle ASM Command-Line Utility

This chapter describes the Oracle Automatic Storage Management (Oracle ASM) Command-Line Utility (ASMCMD). This chapter contains the following topics:

- [About ASMCMD](#)
- [ASMCMD Instance Management Commands](#)
- [ASMCMD File Management Commands](#)
- [ASMCMD Disk Group Management Commands](#)
- [ASMCMD Template Management Commands](#)
- [ASMCMD File Access Control Commands](#)
- [ASMCMD Volume Management Commands](#)

About ASMCMD

ASMCMD is a command-line utility that you can use to manage Oracle ASM instances, disk groups, file access control for disk groups, files and directories within disk groups, templates for disk groups, and volumes.

You can run the ASMCMD utility in either interactive or noninteractive mode.

This section contains the following topics:

- [Types of ASMCMD Commands](#)
- [About Oracle ASM Files, Filenames, Directories, and Aliases](#)
- [Preparing to Run ASMCMD](#)
- [Running ASMCMD in Interactive Mode](#)
- [Running ASMCMD in Noninteractive Mode](#)
- [Getting Help](#)

Types of ASMCMD Commands

The types of ASMCMD commands are listed in [Table 12-1](#).

Table 12–1 Types of ASMCMD Commands

Category	Commands
ASMCMD Instance Management Commands	dsget, dsset, lsct, lsop, lspwusr, orapwusr, shutdown, spbackup, spcopy, spget, spmove, spset, startup
ASMCMD File Management Commands	cd, cp, du, find, ls, lsof, mkalias, pwd, rm, rmalias
ASMCMD Disk Group Management Commands	chdg, chkdg, dropdg, iostat, lsattr, lsdg, lsdsk, lsod, md_backup, md_restore, mkdg, mount, offline, online, rebal, remap, setattr, umount
ASMCMD Template Management Commands	chtmpl, lstmpl, mktmpl, rmtmpl
ASMCMD File Access Control Commands	chgrp, chmod, chown, groups, grpmod, lsgrp, lsusr, mkgrp, mkusr, passwd, rmgrp, rmusr
ASMCMD Volume Management Commands	volcreate, voldelete, voldisable, volenable, volinfo, volresize, volset, volstat

About Oracle ASM Files, Filenames, Directories, and Aliases

ASMCMD works with Oracle ASM files, directories, and aliases. Before using ASMCMD, you should understand how these common computing concepts apply to the Oracle ASM environment.

This section contains the following topics about some key definitions:

- [System-Generated Filename or Fully Qualified Filename](#)
- [Directory](#)
- [Alias](#)
- [Absolute Path and Relative Path](#)
- [Wildcard Characters](#)

System-Generated Filename or Fully Qualified Filename

Every file created in Oracle ASM gets a system-generated filename, otherwise known as a fully qualified filename. This is the same as a complete path name in a local file system.

Oracle ASM generates filenames according to the following scheme:

```
+diskGroupName/databaseName/fileType/fileTypeTag.fileName.incarnation
```

An example of a fully qualified filename is the following:

```
+data/orcl/CONTROLFILE/Current.256.541956473
```

In the previous fully qualified filename, `data` is the disk group name, `orcl` is the database name, `CONTROLFILE` is the file type, and so on.

Only the slash (/) is supported by ASMCMD. Filenames are not case sensitive, but are case retentive. If you type a path name as lowercase, ASMCMD retains the lowercase.

For more information about Oracle ASM filenames, refer to "[Fully qualified File Name Form](#)" on page 7-4.

Directory

As in other file systems, an Oracle ASM directory is a container for files, and an Oracle ASM directory can be part of a tree structure of other directories. The fully qualified filename represents a hierarchy of directories in which the plus sign (+) represent the

root directory. In each disk group, Oracle ASM automatically creates a directory hierarchy that corresponds to the structure of the fully qualified filenames in the disk group. The directories in this hierarchy are known as system-generated directories.

ASMCMD enables you to move up and down in this directory hierarchy with the `cd` (change directory) command. The `ASMCMD ls` (list directory) command lists the contents of the current directory, while the `pwd` command prints the name of the current directory.

When you start ASMCMD, the current directory is set to root (+). For an Oracle ASM instance with two disk groups, for example, `data` and `fra`, entering an `ls` command with the root directory as the current directory produces the following output:

```
ASMCMD> ls
data/
fra/
```

The following example demonstrates navigating the Oracle ASM directory tree (refer to the fully qualified filename shown previously):

```
ASMCMD> cd +data/orcl/CONTROLFILE
ASMCMD> ls
Current.256.541956473
Current.257.541956475
```

You can create your own directories as subdirectories of the system-generated directories using the `ASMCMD mkdir` command. The directories that you create can have subdirectories, and you can navigate the hierarchy of both system-generated directories and user-created directories with the `cd` command.

The following example creates the directory `mydir` under `orcl` in the disk group `data`:

```
ASMCMD> mkdir +data/orcl/mydir
```

Note: The directory `orcl` is a system-generated directory. The contents of `data` represent the contents of disk group `data`.

If you start ASMCMD with the `-p` flag, then ASMCMD shows the current directory as part of its prompt. See ["Including the Current Directory in the ASMCMD Prompt"](#) on page 12-8.

```
ASMCMD [+] > cd data/orcl
ASMCMD [+data/orcl] >
```

ASMCMD retains the case of the directory that you entered.

Alias

Aliases are filenames that are references or pointers to system-generated filenames. However, aliases are user-friendly names. Aliases are similar to symbolic links in UNIX or Linux computers. You can create aliases to simplify Oracle ASM filename administration. You can create aliases with the `mkaalias` ASMCMD command or a SQL `ALTER DISKGROUP` command.

An alias has at a minimum the disk group name as part of its complete path. You can create aliases at the disk group level or in any system-generated or user-created subdirectory. The following are examples of aliases:

```
+data/ctl1.f
+data/orcl/ctl1.f
+data/mydir/ctl1.f
```

If you run the ASMCMD `ls` (list directory) with the `-l` flag, each alias is listed with the system-generated file to which the alias refers.

```
ctl1.f => +data/orcl/CONTROLFILE/Current.256.541956473
```

For more information about aliases, refer to ["Alias Oracle ASM Filename Forms"](#) on page 7-5.

Absolute Path and Relative Path

When you run an ASMCMD command that accepts a filename or directory name as an argument, you can use the name as either an absolute path or a relative path.

An absolute path refers to the full path of a file or directory. An absolute path begins with a plus sign (+) followed by a disk group name, followed by subsequent directories in the directory tree. The absolute path includes directories until the file or directory of interest is reached. A complete system-generated filename, otherwise known as the fully qualified filename, is an example of an absolute path to a file.

Using an absolute path enables the command to access the file or directory regardless of where the current directory is set. The following `rm` command uses an absolute path for the filename:

```
ASMCMD [+] > rm +data/orcl/datafile/users.259.555341963
```

The following `cd` command uses an absolute path to the directory.

```
ASMCMD [+data/mydir] > cd +data/orcl/CONTROLFILE
```

A relative path includes only the part of the filename or directory name that is not part of the current directory. That is, the path to the file or directory is relative to the current directory.

In the following example, the `rm` command operates on the file `undotbs1.272.557429239`, which in this case is a relative path. ASMCMD appends the current directory to the command argument to obtain the absolute path to the file. In this example this is `+data/orcl/DATAFILE/undotbs1.272.557429239`.

```
ASMCMD [+] > cd +data
ASMCMD [+data] > cd orcl/DATAFILE
ASMCMD [+data/orcl/DATAFILE] > ls
EXAMPLE.269.555342243
SYSAUX.257.555341961
SYSTEM.256.555341961
UNDOTBS1.258.555341963
UNDOTBS1.272.557429239
USERS.259.555341963
ASMCMD [+data/orcl/DATAFILE] > rm undotbs1.272.557429239
```

Paths to directories can also be relative. You can go up or down the hierarchy of the current directory tree branch by providing a directory argument to the `cd` command whose path is relative to the current directory.

In addition, you can use the pseudo-directories `"."` and `".."` in place of a directory name. The `"."` pseudo-directory is the current directory. The `".."` pseudo-directory is the parent directory of the current directory.

The following example demonstrates how to use relative directory paths and pseudo-directories:

```
ASMCMD [+data/orcl] > cd DATAFILE
ASMCMD [+data/orcl/DATAFILE] >cd ..
ASMCMD [+data/orcl] >
```

Wildcard Characters

The wildcard characters * and % match zero or more characters anywhere within an absolute or relative path, which saves typing of the full directory or file name. The two wildcard characters behave identically. There are various ASMCMD commands that accept wildcards. These include `cd`, `du`, `find`, `ls`, `lsattr`, `lsdgs`, `lsdsk`, `lsgrp`, `lsusr`, and `rm`.

If a wildcard pattern matches only one directory when using wildcard characters with `cd`, then `cd` changes the directory to that destination. If the wildcard pattern matches multiple directories, then ASMCMD does not change the directory but instead returns an error.

[Example 12-1](#) illustrates the use of wildcards.

Example 12-1 Using Wildcards With ASMCMD

```
ASMCMD [+] > cd +data/orcl/*FILE
ASMCMD-08005: +data/orcl/*FILE: ambiguous

ASMCMD [+] > cd +data/orcl/C*
ASMCMD [+data/orcl/CONTROLFILE] >

ASMCMD [+] > ls +fra/orcl/A%
2009_07_13/
2009_07_14/

ASMCMD [+] > ls +fra/orcl/ARCHIVELOG/2009%

+fra/orcl/ARCHIVELOG/2009_07_13/:
thread_1_seq_3.260.692103543
thread_1_seq_4.261.692108897
thread_1_seq_5.262.692125993
thread_1_seq_6.263.692140729
thread_1_seq_7.264.692143333

+fra/orcl/ARCHIVELOG/2009_07_14/:
thread_1_seq_8.271.692158265
thread_1_seq_9.272.692174597

ASMCMD [+] > ls data/orcl/*

+data/orcl/CONTROLFILE/:
Current.260.692103157

+data/orcl/DATAFILE/:
EXAMPLE.265.692103187
SYSaux.257.692103045
SYSTEM.256.692103045
UNDOTBS1.258.692103045
USERS.259.692103045

+data/orcl/ONLINELOG/:
group_1.261.692103161
```

```
group_2.262.692103165
group_3.263.692103169

+data/orcl/PARAMETERFILE/:
spfile.266.692103315

+data/orcl/TEMPFILE/:
TEMP.264.692103181
spfileorcl.ora
```

Preparing to Run ASMCMD

You can run the ASMCMD utility in either interactive or noninteractive mode.

Before running ASMCMD, review the items in the following list.

- Log in to the host on which the Oracle ASM instance that you want to work with is running.

You must log in as a user that has `SYSASM` or `SYSDBA` privileges through operating system authentication. The `SYSASM` privilege is the required connection to administer the Oracle ASM instance. See "[Authentication for Accessing Oracle ASM Instances](#)" on page 3-22.

See Also: Refer to the *Oracle Database Administrator's Guide* for information about operating system authentication

- To connect to the Oracle ASM instance, run ASMCMD that is located in the `bin` subdirectory of the Oracle grid infrastructure home (Oracle ASM home).

Connect as `SYSASM`, the default connection, to administer an Oracle ASM instance.

Ensure that the `ORACLE_HOME` and `ORACLE_SID` environment variables refer to the Oracle ASM instance. Depending on your operating system, you might have to set other environment variables to properly connect to the Oracle ASM instance.

Ensure that the `bin` subdirectory of your Oracle grid infrastructure home is in your `PATH` environment variable.

See Also: Refer to the *Oracle Database Administrator's Guide* for more information about setting environment variables

The default value of the Oracle ASM SID for a single-instance database is `+ASM`. In Real Application Clusters environments, the default value of the Oracle ASM SID on any node is `+ASMnode#`.

- To use most of the ASMCMD commands, ensure that the Oracle ASM instance is started and the Oracle ASM disk groups are mounted.

If the Oracle ASM instance is not running or if the `ORACLE_SID` is set incorrectly, ASMCMD runs only those commands that do not require an Oracle ASM instance. The commands include `startup`, `shutdown`, `lsdisk`, `help`, and `exit`. If you attempt to run other ASMCMD commands, an error message displays.

- You can connect to the database instance as `SYSDBA` by running ASMCMD that is located in the `bin` directory of the Oracle Database home.

Ensure that the `ORACLE_HOME` and `ORACLE_SID` environment variables refer to the database instance. Depending on your operating system, you might have to set other environment variables to properly connect to the database instance.

You must include the `-a` option to connect as SYSDBA. See ["Specifying the Type of Connection"](#) on page 12-8.

With this connection, there are a limited set of operations that can be run. For more information, see ["The SYSDBA Privilege for Managing Oracle ASM Components"](#) on page 3-26.

When administering disk groups, Oracle recommends that you run ASMCMD from the database home of the database instance that is the owner of the files in the disk group.

Running ASMCMD in Interactive Mode

The interactive mode of the ASMCMD utility provides a shell-like environment where you are prompted to enter ASMCMD commands. The syntax for starting ASMCMD in interactive mode is:

```
asmcmd [-v]
asmcmd [-v] [-a connection_type] [-p]
```

[Table 12-2](#) summarizes the ASMCMD options.

Table 12-2 ASMCMD Interactive Options

Option	Description
<code>-v</code>	Displays the version of ASMCMD and then exits.
<code>-v</code>	Displays additional information with some commands to help users diagnose problems.
<code>-a <i>connection_type</i></code>	Specifies how to connect as when accessing an Oracle ASM or database instance. Can be either SYSASM or SYSDBA. The default is SYSASM and should be used when administering the Oracle ASM instance.
<code>-p</code>	Displays the current directory in the prompt.

ASMCMD can be started with multiple options. For example:

```
asmcmd -p -v
```

For the majority of the examples in this chapter, ASMCMD was started with the `-p` and `-v` options

Running ASMCMD Without Any Options

To run ASMCMD in interactive mode without any options:

1. Enter the following at the operating system command prompt:

```
asmcmd
```

Oracle displays an ASMCMD command prompt as follows:

```
ASMCMD>
```

2. Enter an ASMCMD command and press `Enter`. The command runs and displays its output, if any, and then ASMCMD prompts for the next command.
3. Continue entering ASMCMD commands until you have completed the tasks.
4. Enter the `exit` or `quit` command to exit ASMCMD.

Specifying the Type of Connection

You can specify the `-a` option to choose the type of connection, either `SYSASM` or `SYSDBA`. The default value is `SYSASM` and should be used when administering an Oracle ASM instance. Connect as `SYSDBA` when connecting to the database instance.

For example:

```
$ asmcmd -a sysasm
```

Specifying the Verbose Mode

You can specify the `-v` option with the `asmcmd` command to display additional information with some commands to help users diagnose problems as shown in the following example:

```
$ asmcmd -v
```

Including the Current Directory in the ASMCMD Prompt

You can specify the `-p` option with the `asmcmd` command to include the current directory in the ASMCMD prompt as shown in the following example:

```
$ asmcmd -p
ASMCMD [+] > cd data
ASMCMD [+data] >
```

Displaying the ASMCMD Version Number

You can specify the `-V` option when starting `asmcmd` to displays the `asmcmd` version number. After displaying the version number, `asmcmd` immediately exits.

For example:

```
$ asmcmd -V
asmcmd version 11.2.0.1.0
$
```

Running ASMCMD in Noninteractive Mode

In noninteractive mode, you run a single ASMCMD command by including the command and command options on the command line that invokes ASMCMD. ASMCMD runs the command, generates output if any, and then exits. The noninteractive mode is especially useful for running scripts.

To run ASMCMD in noninteractive mode where *command* is any valid ASMCMD command and *options* is a list of command options, at the command prompt enter the following:

```
asmcmd command options
```

[Example 12-2](#) shows how to run ASMCMD in the noninteractive mode. The first example uses the `ls` command to list the disk groups for the Oracle ASM instance. The second example redirects the output of the `lsod` command to the `my_lsod_test` file.

Example 12-2 Running ASMCMD in Noninteractive Mode

```
$ asmcmd ls -l
State   Type   Rebal  Name
MOUNTED NORMAL N      DATA/
MOUNTED NORMAL N      FRA/

$ asmcmd lsod -H -G data --process LGWR > my_lsod_test
```


Getting Help

Type `help` at the ASMCMD prompt or as a command in noninteractive mode to view general information about ASMCMD and a list of available ASMCMD commands.

You can type `help command` to display help text for a specific command, including usage information about how to run the command with its options.

The following is an example of the use of the `help` command.

Example 12–3 Displaying ASMCMD Help Text

```
ASMCMD [+] > help startup
startup [--nomount] [--restrict] [--pfile <pfile.ora>]

Start the ASM instance.

[--nomount]           specifies the nomount option.
[--restrict]          start the instance in restricted mode.
[--pfile <pfile.ora>] specifies the location of the pfile.
```

ASMCMD Instance Management Commands

This section describes the ASMCMD Oracle ASM instance management commands.

Some commands in this section affect the Grid Plug and Play (GPnP) profile, which is a resource in a clustered configuration. In an Oracle Restart configuration, the profile is actually located in a resource attribute, not the Grid Plug and Play (GPnP) profile. However, the functionality of the commands is the same for both configurations.

[Table 12–3](#) provides a summary of the Oracle ASM instance management commands.

Table 12–3 Summary of ASMCMD Instance Management Commands

Command	Description
<code>dsget</code>	Retrieves the discovery diskstring value that is used by the Oracle ASM instance and its clients.
<code>dsset</code>	Sets the disk discovery diskstring value that is used by the Oracle ASM instance and its clients.
<code>lsct</code>	Lists information about current Oracle ASM clients.
<code>lsop</code>	Lists the current operations on a disk group or Oracle ASM instance.
<code>lspwusr</code>	List the users from an Oracle ASM password file.
<code>orapwusr</code>	Add, drop, or change an Oracle ASM password user.
<code>shutdown</code>	Shuts down an Oracle ASM instance.
<code>spbackup</code>	Backs up an Oracle ASM SPFILE.
<code>spcopy</code>	Copies an Oracle ASM SPFILE.
<code>spget</code>	Retrieves the location of the Oracle ASM SPFILE.
<code>spmmove</code>	Moves an Oracle ASM SPFILE.
<code>spset</code>	Sets the location of the Oracle ASM SPFILE.
<code>startup</code>	Starts up an Oracle ASM instance.

dsget

Purpose

Retrieves the discovery diskstring value that is used by the Oracle ASM instance and its clients.

Syntax and Description

```
dsget [[--normal] [--profile [--force]] [--parameter]]
```

The syntax options for the dsget command are described in [Table 12–4](#).

Table 12–4 Options for the dsget command

Option	Description
--normal	Retrieves the discovery string from the Grid Plug and Play (GPnP) profile and the one that is set in the Oracle ASM instance. It returns one row each for the profile and parameter setting. This is the default setting.
--profile [--force]	Retrieves the discovery string from the GPnP profile. If --force is specified, retrieves the discovery string from the local GPnP profile.
--parameter	Retrieves the ASM_DISKSTRING parameter setting of the Oracle ASM instance.

Example

The following example uses dsget to retrieve the current discovery diskstring value from the GPnP profile and the ASM_DISKSTRING parameter.

Example 12–4 Using dsget

```
ASMCMD [+] > dsget
profile: /devices/disk*
parameter: /devices/disk*
```

dsset

Purpose

Sets the discovery diskstring value that is used by the Oracle ASM instance and its clients.

The specified diskstring must be valid for existing mounted disk groups. The updated value takes effect immediately.

Syntax and Description

```
dsset [--normal] [--parameter] [--profile [--force]] diskstring
```

The syntax options for the dsset command are described in [Table 12–5](#).

Table 12–5 Options for the dsset command

Option	Description
--normal	Sets the discovery string in the Grid Plug and Play (GPnP) profile and in the Oracle ASM instance. The update occurs after the Oracle ASM instance has successfully validated that the specified discovery string has discovered all the necessary disk groups and voting files. This command fails if the instance is not using a server parameter file (SPFILE). This is the default setting.
--profile [--force]	Specifies the discovery diskstring that is pushed to the GPnP profile without any validation by the Oracle ASM instance, ensuring that the instance can discover all the required disk groups. The update is guaranteed to be propagated to all the nodes that are part of the cluster. If --force is specified, the specified diskstring is pushed to the local GPnP profile without any synchronization with other nodes in the cluster. This command option updates only the local profile file. This option should only be used in the case of recovery. The command fails if the Oracle Clusterware stack is running.
--parameter	Specifies that the diskstring is updated in memory after validating that the discovery diskstring discovers all the current mounted disk groups and voting files. The diskstring is not persistently recorded in either the SPFILE or the GPnP profile.
<i>diskstring</i>	Specifies the value for the discovery diskstring.

For information about disk discovery and the discovery diskstring, see "[Oracle ASM Disk Discovery](#)" on page 4-18.

Example

The following example uses dsset to set the current value of the discovery diskstring in the GPnP profile.

Example 12–5 Using dsset

```
ASMCMD [+] > dsset /devices/disk*
```

lsct

Purpose

Lists information about current Oracle ASM clients from the V\$ASM_CLIENT view. A client, such as Oracle Database or Oracle ASM Dynamic Volume Manager (Oracle ADVM), uses disk groups that are managed by the Oracle ASM instance to which ASMCMD is currently connected.

Syntax and Description

```
lsct [-gH] [diskgroup]
```

[Table 12–6](#) lists the options for the lsct command.

Table 12–6 Options for the lsct command

Option	Description
(none)	Displays information about current clients from the V\$ASM_CLIENT view.

Table 12–6 (Cont.) Options for the lsct command

Option	Description
-g	Selects from the GV\$ASM_CLIENT view. GV\$ASM_CLIENT.INST_ID is included in the output.
-H	Suppresses column headings.
<i>diskgroup</i>	Specifies the disk group.

If *diskgroup* is specified, then only information about that disk group is listed.

Example

The following example displays information about the clients that are accessing the data disk group.

Example 12–6 Using lsct

```
ASMCMD [+] > lsct data
DB_Name      Status      Software_Version  Compatible_version  Instance_Name  Disk_Group
+ASM         CONNECTED   11.2.0.1.0       11.2.0.0.0         +ASM           DATA
asmvol       CONNECTED   11.2.0.1.0       11.2.0.0.0         +ASM           DATA
orcl         CONNECTED   11.2.0.1.0       11.2.0.0.0         orcl           DATA
```

lsop

Purpose

Lists the current operations on a disk group or Oracle ASM instance.

Syntax and Description

```
lsop
```

lsop displays information from the V\$ASM_OPERATION view.

Example

The following are examples of the lsop command. The examples list operations on the disk groups of the current Oracle ASM instance.

Example 12–7 Using lsop

```
ASMCMD [+] > lsop
Group_Name  Dsk_Num  State  Power
DATA        REBAL    WAIT   2

ASMCMD [+] > lsop
Group_Name  Dsk_Num  State  Power
FRA         REBAL    REAP   3
```

lspwusr

Purpose

List the users from the local Oracle ASM password file.

Syntax and Description

```
lspwusr [-H]
```

The options for the `lspwusr` command are as follows:

Table 12-7 Options for the `lspwusr` Command

Option	Description
-H	Suppresses column headers from the output.

Examples

The following is an example of the `lspwusr` example. The example lists the current current users in the local Oracle ASM password file.

Example 12-8 Using `lspwusr`

```
ASMCMD [+] > lspwusr
Username sysdba sysoper sysasm
        SYS  TRUE   TRUE   TRUE
ASMSNMP  TRUE  FALSE  FALSE
```

orapwusr

Purpose

Add, drop, or modify an Oracle ASM password file user.

Syntax and Description

```
orapwusr { { { --add | --modify [--password] }
  [--privilege {sysasm|sysdba|sysoper} ] } | --delete } user
```

The options for the `orapwusr` command are as follows:

Table 12-8 Options for the `orapwusr` Command

Option	Description
--add	Adds a user to the password file. Also prompts for a password.
--delete	Drops a user from the password file.
--modify	Changes a user in the password file.
--privilege <i>role</i>	Sets the role for the user. The options are <code>sysasm</code> , <code>sysdba</code> , and <code>sysoper</code> .
--password	Prompts for and then changes the password of a user.
<i>user</i>	Name of the user to add, drop, or modify.

`orapwusr` attempts to update passwords on all nodes in a cluster. The command requires the `SYSASM` privilege to run. A user logged in as `SYSDBA` cannot change its password using this command.

Examples

The following is an example of the `orapwusr` command. This example adds the `hruser` to the Oracle ASM password file with the role of the user set to `SYSDBA`.

Example 12-9 Using `orapwusr`

```
ASMCMD [+] > orapwusr --add --privilege sysdba hruser
```

shutdown

Purpose

Shuts down an Oracle ASM instance.

Syntax and Description

shutdown [--abort | --immediate]

Table 12–9 Options for the shutdown command

Option	Description
--abort	Shut down aborting all existing operations.
--immediate	Shut down immediately.

The default action is a normal shutdown.

Oracle strongly recommends that you shut down all database instances that use the Oracle ASM instance and dismount all file systems mounted on Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes before attempting to shut down the Oracle ASM instance with the abort (--abort) option.

For more information about shutting down an Oracle ASM instance, see ["Shutting Down an Oracle ASM Instance"](#) on page 3-16.

Example

The following are examples of the shutdown command. The first example performs a shut down of the Oracle ASM instance with normal action. The second example performs a shut down with immediate action. The third example performs a shut down that aborts all existing operations.

Example 12–10 Using shutdown

```
ASMCMD [+] > shutdown
ASMCMD [+] > shutdown --immediate
ASMCMD [+] > shutdown --abort
```

spbackup

Purpose

Backs up an Oracle ASM SPFILE.

Syntax and Description

spbackup *source destination*

Table 12–10 Options for the spbackup command

Option	Description
<i>source</i>	Specifies the source file name.
<i>destination</i>	Specifies the destination file.

`spbackup` does not affect the GPnP profile. The backup file that is created is not a special file type and is not identified as a SPFILE. This file cannot be copied with `spcopy`. To copy this backup file, use the ASMCMD `cp` command.

Example

The following are examples of the `spbackup` command. The first example backs up the Oracle ASM SPFILE from one operating system location to another. The second example backs up the SPFILE from an operating system location to the `data/bakspfileASM.ora` disk group.

Example 12–11 Using `spbackup`

```
ASMCMD> spbackup /u01/oracle/dbs/spfile+ASM.ora /u01/oracle/dbs/bakspfileASM.ora
ASMCMD> spbackup /u01/oracle/dbs/spfile+ASM.ora +DATA/bakspfileASM.ora
```

spcopy

Purpose

Copies an Oracle ASM SPFILE from source to destination.

Syntax and Description

```
spcopy [-u] source destination
```

Table 12–11 Options for the `spcopy` command

Option	Description
<code>-u</code>	Updates the Grid Plug and Play (GPnP) profile.
<code>source</code>	Specifies the source file name.
<code>destination</code>	Specifies the destination.

Note the following about the use of `spcopy`:

- `spcopy` can copy an Oracle ASM SPFILE from a disk group to a different disk group or to an operating system file.
- `spcopy` can copy an Oracle ASM SPFILE from an operating system file to a disk group or to an operating system file.
- `spcopy` can copy an Oracle ASM SPFILE when the SPFILE is being used by an open Oracle ASM instance.

Include the `-u` option to update the GPnP profile. You can also use `spset` to update the GPnP profile. See "[spset](#)" on page 12-17. For information about copying and moving an Oracle ASM instance initialization parameter file after upgrading, see "[Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File](#)" on page 3-4.

After copying the SPFILE, you must restart the instance with the SPFILE in the new location if you want to use that SPFILE. When the Oracle ASM instance is running with the SPFILE in the new location, you can remove the source SPFILE.

To use `spcopy` to copy an Oracle ASM SPFILE into a disk group, the disk group attribute `COMPATIBLE.ASM` must be set to 11.2 or greater.

Example

The following are examples of the `spcopy` command. The first example copies the Oracle ASM SPFILE from one operating system location to another. The second example copies the SPFILE from an operating system location to the `data` disk group and updates the GPnP profile.

Example 12–12 Using `spcopy`

```
ASMCMD> spcopy /u01/oracle/dbs/spfile+ASM.ora /u01/oracle/dbs/testspfileASM.ora
```

```
ASMCMD> spcopy -u /u01/oracle/dbs/spfile+ASM.ora +DATA/testspfileASM.ora
```

See Also: The `CREATE SPFILE` SQL statement in the *Oracle Database SQL Language Reference* for information about creating a server parameter file

spget

Purpose

Retrieves the location of the Oracle ASM SPFILE from the Grid Plug and Play (GPnP) profile.

Syntax and Description

```
spget
```

The location retrieved by `spget` is the location in the GPnP profile, but not always the location of the SPFILE currently used. For example, the location could have been recently updated by `spset` or `spcopy` with the `-u` option on an Oracle ASM instance that has not been restarted. After the next restart of the Oracle ASM, this location points to the ASM SPFILE currently being used.

Example

The following is an example of the `spget` command that retrieves and displays the location of the SPFILE from the GPnP profile.

Example 12–13 Using `spget`

```
ASMCMD [+] > spget  
+DATA/asm/asmparameterfile/registry.253.691575633
```

spmove

Purpose

Moves an Oracle ASM SPFILE from source to destination and automatically updates the GPnP profile.

Syntax and Description

```
spmove source destination
```


Table 12–12 Options for the *spmove* command

Option	Description
<i>source</i>	Specifies the source file name.
<i>destination</i>	Specifies the destination file.

Note the following about the use of *spmove*:

- *spmove* can move an Oracle ASM SPFILE when the open instance is using a PFILE or a different SPFILE. After moving the SPFILE, you must restart the instance with the SPFILE in the new location if you want to use that SPFILE.
- *spmove* cannot move an Oracle ASM SPFILE when the SPFILE is being used by an open Oracle ASM instance.

For information about copying and moving an Oracle ASM instance initialization parameter file after upgrading, see "[Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File](#)" on page 3-4.

To use *spmove* to move an Oracle ASM SPFILE into a disk group, the disk group attribute `COMPATIBLE.ASM` must be set to 11.2 or greater.

Example

The following are examples of the *spmove* command. The first example moves the Oracle ASM SPFILE from one operating system location to another. The second example moves the SPFILE from an operating system location to the data disk group.

Example 12–14 Using *spmove*

```
ASMCMD> spmove /u01/oracle/dbs/spfile+ASM.ora /u01/oracle/dbs/testspfileASM.ora
ASMCMD> spmove /u01/oracle/dbs/spfile+ASM.ora +DATA/testspfileASM.ora
```

spset

Purpose

Sets the location of the Oracle ASM SPFILE in the Grid Plug and Play (GPnP) profile.

Syntax and Description

spset location

Table 12–13 Options for the *spset* command

Option	Description
<i>diskgroup</i>	Specifies the disk group that is the location of the Oracle ASM SPFILE.
<i>location</i>	Specifies the location of the Oracle ASM SPFILE. The location is the path to an Oracle ASM SPFILE.

Example

The following is an example of the *spset* command that sets the location of the Oracle ASM SPFILE command in the data disk group.

Example 12–15 Using spset

```
ASMCMD> spset +DATA//asm/asmparameterfile/asmspfile.ora
```

startup

Purpose

Starts up an Oracle ASM instance.

Syntax and Description

```
startup [--nomount][--restrict][--pfile pfile]
```

Table 12–14 Options for the startup command

Option	Description
--nomount	Specifies no mount operation.
--restrict	Specifies restricted mode.
--pfile <i>pfile</i>	Oracle ASM initialization parameter file.

The default action is a startup that mounts disk groups and enables Oracle ADVM volumes. For information about disk groups that are mounted at startup time, see ["About Mounting Disk Groups at Startup"](#) on page 3-15.

For more information about starting up an Oracle ASM instance, see ["Starting Up an Oracle ASM Instance"](#) on page 3-13.

Example

The following is an example of the `startup` command that starts the Oracle ASM instance without mounting disk groups and uses the `asm_init.ora` initialization parameter file.

Example 12–16 Using startup

```
ASMCMD> startup --nomount --pfile asm_init.ora
```

ASMCMD File Management Commands

This section describes the ASMCMD file management commands.

[Table 12–15](#) provides a summary of the file management commands.

Table 12–15 Summary of ASMCMD File Management Commands

Command	Description
<code>cd</code>	Changes the current directory to the specified directory.
<code>cp</code>	Enables you to copy files between disk groups on a local instance and remote instances.
<code>du</code>	Displays the total disk space occupied by files in the specified Oracle ASM directory and all of its subdirectories, recursively.
<code>find</code>	Lists the paths of all occurrences of the specified name (with wildcards) under the specified directory.

Table 12–15 (Cont.) Summary of ASMCMD File Management Commands

Command	Description
ls	Lists the contents of an Oracle ASM directory, the attributes of the specified file, or the names and attributes of all disk groups.
lsof	Lists the open files.
mkalias	Creates an alias for system-generated filenames.
mkdir	Creates Oracle ASM directories.
pwd	Displays the path of the current Oracle ASM directory.
rm	Deletes the specified Oracle ASM files or directories.
rmbalias	Deletes the specified alias, retaining the file that the alias points to.

cd

Purpose

Changes the current directory to the specified directory.

Syntax and Description

`cd [dir]`

Table 12–16 Options for the cd Command

Option	Description
<i>dir</i>	Name of the directory.

dir can be specified as either an absolute path or a relative path, including the `.` and `..` pseudo-directories. *dir* can contain wildcard characters. See "[Wildcard Characters](#)" on page 12-5.

Examples

The following are examples of the `cd` command changing into various directories.

Example 12–17 Using cd

```
ASMCMD [+data/hr] > cd +data/orcl
```

```
ASMCMD [+data/orcl] > cd DATAFILE
```

```
ASMCMD [+data/orcl/DATAFILE] >cd ..
```

cp

Purpose

Enables you to copy files between Oracle ASM disk groups on local instances to and from remote instances.

Syntax and Description

`cp [-i] [-f] [connect_str:]src_file [connect_str:]tgt_file`

The syntax options for the `cp` command are described in [Table 12–17](#).

Table 12–17 Options for the cp Command

Option	Description
-i	Interactive, prompt before copy file or overwrite
-f	Force, if an existing destination file, remove it and try again without user interaction
<i>connect_str</i>	The connection string for use with a remote instance copy.
<i>src_file</i>	Name of the source file to copy.
<i>tgt_file</i>	A user alias for the created target file name or alias directory name.

cp cannot copy files between two remote instances. The local Oracle ASM instance must be either the source or the target of the operation.

You can use the cp command to:

- Copy files from a disk group to the operating system
- Copy files from a disk group to a disk group
- Copy files from the operating system to a disk group

Some file types cannot be the source or destination of the cp command. These file types include OCR and OCR backup file types. To back up, copy, or move an ASM SPFILE, use the spbackup, spcopy, or spmove commands.

connect_str is not required for a local instance copy, which is the default case. In the case of a remote instance copy, you must specify the connect string and Oracle ASM prompts for a password in a non-echoing prompt. The *connect_str* is in the form of:

```
user@host[.port_number].SID
```

user, *host*, and *SID* are required in the *connect_str* parameter. The default port number is 1521.

See Also: *Oracle Database Net Services Administrator's Guide* for more information about connect strings

src_file must be either the fully qualified file name, the system-generated name, or the Oracle ASM alias.

The cp command performs a bit-wise copy. There is no data transformation.

Examples

The following are examples of the cp command. The first example shows a copy of a file in the data disk group to a file on the operating system. The second example shows a copy of a file on the operating system to the data disk group.

Example 12–18 Using cp

```
ASMCMD [+] > cp +data/orcl/datafile/EXAMPLE.265.691577295 /mybackups/example.bak
copying +data/orcl/datafile/EXAMPLE.265.691577295 -> /mybackups/example.bak
```

```
ASMCMD [+] > cp /mybackups/examples.bak +data/orcl/datafile/myexamples.bak
copying /mybackups/examples.bak -> +data/orcl/datafile/myexamples.bak
```

du

Purpose

Displays the total space used for files in the specified directory and in the entire directory tree under the directory.

Syntax and Description

```
du [-H] [dir]
```

Table 12-18 Options for the du command

Option	Description
<i>dir</i>	Name of the directory.
-H	Suppresses column headings from the output.

If you do not specify *dir*, then information about the current directory is displayed. *dir* can contain wildcard characters. See "[Wildcard Characters](#)" on page 12-5.

The following two values are displayed, both in units of megabytes.

- `Used_MB` - This value does not include mirroring.
- `Mirror_used_MB` - This value includes mirroring.

For example, if a normal redundancy disk group contains 100 MB of data, then assuming that each file in the disk group is 2-way mirrored, `Used_MB` is 100 MB and `Mirror_used_MB` is roughly 200 MB.

Example

The following is an example of the `du` command. The example shows disk space used in the `orcl` directory in `data`, including all of the directories under the `orcl` directory.

Example 12-19 Using du

```
ASMCMD [+] > du data/orcl
Used_MB      Mirror_used_MB
  1756             3519
```

find

Purpose

Displays the absolute paths of all occurrences of the specified name pattern (with wildcards) in a specified directory and its subdirectories.

Syntax and Description

```
find [--type type] dir pattern
```

Table 12-19 Options for the find Command

Option	Description
--type <i>type</i>	Type of target that you want to find.
<i>dir</i>	Directory name where you want to start searching.
<i>pattern</i>	Name of a target or a wildcard pattern.

This command searches the specified directory and all subdirectories under it in the directory tree for the supplied *pattern*. The value that you use for *pattern* can be a directory name or a filename, and can include wildcard characters. See ["Wildcard Characters"](#) on page 12-5.

The ASMCMD `find` command is case insensitive.

In the output of the command, directory names are suffixed with the slash character (/) to distinguish them from filenames.

You use the `--type` flag to find all the files of a particular type (specified as *type*). For example, you can search for control files by specifying *type* as `CONTROLFILE`. Valid values for *type* are listed in [Table 7-1, "File Types Supported by Automatic Storage Management"](#) on page 7-1. These are type values from the *type* column of the `V$ASM_FILE` view.

Examples

The following are examples of the `find` command. The first example searches the data disk group for files that begin with UNDO. The second example returns the absolute path of all the control files (`--type CONTROLFILE`) in the `+data/orcl` directory.

Example 12-20 Using find

```
ASMCMD [+] > find +data undo*
+data/ORCL/DATAFILE/UNDOTBS1.258.691577151

ASMCMD [+] > find --type CONTROLFILE +data/orcl *
+data/orcl/CONTROLFILE/Current.260.691577263
```

ls

Purpose

Lists the contents of an Oracle ASM directory, the attributes of the specified file, or the names and attributes of all disk groups.

Syntax and Description

```
ls [-lsdtLagH] [--reverse] [--permission] [pattern]
```

Table 12-20 Options for the ls command

Option	Description
(none)	Displays only filenames and directory names.
-l	Displays extended file information, including striping and redundancy information and whether the file was system-generated (indicated by Y under the SYS column) or user-created (as in the case of an alias, indicated by N under the SYS column). When used in the "ls -l +" command, displays directory information. Not all possible file attributes or disk group attributes are included. To view the complete set of column values for a file or a disk group, query the <code>V\$ASM_FILE</code> and <code>V\$ASM_DISKGROUP</code> views.
-s	Displays file space information.
-d	If the value for the <i>name</i> argument is a directory, then <code>ls</code> displays information about that directory, rather than the directory contents. Typically used with another flag, such as the <code>-l</code> flag.

Table 12–20 (Cont.) Options for the ls command

Option	Description
--reverse	Reverses the sort order of the listing.
-t	Sorts the listing by timestamp (latest first) instead of by name.
-L	If the value for the <i>name</i> argument is an alias, then ASMCMD displays information about the file that it references. Typically used with another flag, such as the -l flag.
-a	For each listed file, displays the absolute path of the alias that references it, if any.
-g	GV\$ASM_DISKGROUP.INST_ID is included in the output.
-H	Suppresses column headings.
--permission	Shows the permissions of a file (V\$ASM_FILE.permission, V\$ASM_FILE.owner, V\$ASM_FILE.usergroup, V\$ASM_ALIAS.name).
<i>pattern</i>	Name of a file, directory, or pattern.

Command options enable you to modify and customize the output of the command. Table 12–20 lists the options and their descriptions. For disk group information, this command queries the V\$ASM_DISKGROUP_STAT view by default. If you specify all of the options, then the command shows a union of their attributes, with duplicates removed. If you enter `ls +`, the top level directory structure is displayed.

name can be a filename or directory name, including wildcard characters. See "Wildcard Characters" on page 12-5.

- If *name* is a directory name, then `ls` lists the contents of the directory and depending on flag settings, `ls` also lists information about each directory member. Directories are listed with a trailing slash (/) to distinguish them from files.
- If the value that you enter for *name* is a filename, then `ls` lists the file and depending on the flag settings, `ls` also lists information about the file. The file must be located in the current directory if the filename is specified with a relative path.

Examples

The following are examples of the `ls` command that display various information about directories and the contents of the directories.

Example 12–21 Using ls

```
ASMCMD [+] > ls +data/orcl/datafile
EXAMPLE.265.691577295
SYS_AUX.257.691577149
SYSTEM.256.691577149
UNDOTBS1.258.691577151
USERS.259.691577151
```

```
ASMCMD [+] > ls -lt +data/orcl/datafile
Type      Redund  Striped  Time                Sys  Name
DATAFILE  MIRROR  COARSE   JUL 13 08:00:00    Y    EXAMPLE.265.691577295
DATAFILE  MIRROR  COARSE   JUL 13 05:00:00    Y    SYS_AUX.257.691577149
DATAFILE  MIRROR  COARSE   JUL 13 02:00:00    Y    USERS.259.691577151
DATAFILE  MIRROR  COARSE   JUL 13 02:00:00    Y    UNDOTBS1.258.691577151
DATAFILE  MIRROR  COARSE   JUL 13 02:00:00    Y    SYSTEM.256.691577149
```

```
ASMCMD [+] > ls -l +data/orcl/datafile/sy*
```

```
Type      Redund  Striped  Time          Sys  Name
DATAFILE  MIRROR  COARSE   JUL 13 05:00:00 Y    SYSAUX.257.691577149
DATAFILE  MIRROR  COARSE   JUL 13 02:00:00 Y    SYSTEM.256.691577149
```

```
ASMCMD [+] > ls -s +data/orcl/datafile
Block_Size  Blocks      Bytes      Space  Name
8192        12801      104865792  214958080 EXAMPLE.265.691577295
8192        88321      723525632  1452277760 SYSAUX.257.691577149
8192        88321      723525632  1452277760 SYSTEM.256.691577149
8192        7681       62922752   131072000 UNDOTBS1.258.691577151
8192        641        5251072    12582912  USERS.259.691577151
```

```
ASMCMD [+] > ls --permission +data/orcl/datafile
User  Group  Permission  Name
      rw-rw-rw-  EXAMPLE.265.691577295
      rw-rw-rw-  SYSAUX.257.691577149
      rw-rw-rw-  SYSTEM.256.691577149
      rw-rw-rw-  UNDOTBS1.258.691577151
      rw-rw-rw-  USERS.259.691577151
```

lsf

Purpose

Lists the open files of the local clients.

Syntax and Description

```
lsf [-H] [-G diskgroup] [--dbname db] [-C instance]
```

Table 12–21 Options for the lsf command

Option	Description
-H	Suppresses column headings.
-G <i>diskgroup</i>	List files only from this specified disk group.
--dbname <i>db</i>	List files only from this specified database.
-C <i>instance</i>	List files only from this specified instance.

Example

The following are examples of the `lsf` command. The first example lists the open files for the data disk group. The second example lists the open files for the Oracle ASM instance.

Example 12–22 Using lsf

```
ASMCMD [+] > lsf -G data
DB_Name  Instance_Name  Path
orcl     orcl           +data/orcl/controlfile/current.260.691577263
orcl     orcl           +data/orcl/datafile/example.265.691577295
orcl     orcl           +data/orcl/datafile/sysaux.257.691577149
orcl     orcl           +data/orcl/datafile/system.256.691577149
orcl     orcl           +data/orcl/datafile/undotbs1.258.691577151
orcl     orcl           +data/orcl/datafile/users.259.691577151
orcl     orcl           +data/orcl/onlinelog/group_1.261.691577267
orcl     orcl           +data/orcl/onlinelog/group_2.262.691577271
orcl     orcl           +data/orcl/onlinelog/group_3.263.691577275
```



```

orcl      orcl      +data/orcl/tempfile/temp.264.691577287

ASMCMD [+] > lsof -C +ASM
DB_Name  Instance_Name Path
asmvol   +ASM      +data/VOLUME1.271.679226013
asmvol   +ASM      +data/VOLUME2.272.679227351
    
```

mkalias

Purpose

Creates an alias for the specified system-generated filename.

Syntax and Description

`mkalias file alias`

Table 12–22 Options for the mkalias command

Option	Description
<i>file</i>	System-generated file name.
<i>alias</i>	Alias for the file name.

alias must be in the same disk group as the system-generated file. Only one alias is permitted for each Oracle ASM file.

Example

The following example creates the `sysaux.f` alias for the fully qualified filename `+data/orcl/DATAFILE/SYSAUX.257.691577149`. Following the `mkalias` command, `ls -a` is run to check the results.

Example 12–23 Using mkalias

```

ASMCMD [+data/orcl/datafile] > mkalias SYSAUX.257.691577149 sysaux.f

ASMCMD [+data/orcl/datafile] > ls -a
none => EXAMPLE.265.691577295
none => SYSTEM.256.691577149
none => UNDOTBS1.258.691577151
none => USERS.259.691577151
+DATA/ORCL/DATAFILE/sysaux.f => SYSAUX.257.691577149
sysaux.f
    
```

mkdir

Purpose

Creates Oracle ASM directories under the current directory.

Syntax and Description

`mkdir dir [dir...]`

Table 12–23 Options for the mkdir command

Option	Description
<i>dir</i>	Directory name that you want to create.

The current directory can be created by the system or by the user. You cannot create a directory at the root (+) level.

Example

The following is an example of the `mkdir` command. The example creates the directories `subdir1` and `subdir2` at the disk group level in the disk group `data`.

Example 12–24 Using mkdir

```
ASMCMD [+data] > mkdir subdir1 subdir2

ASMCMD [+data] > ls
ASM/
ORCL/
subdir1/
subdir2/
```

pwd

Purpose

Displays the absolute path of the current directory.

Syntax and Description

```
pwd
```

Example

The following is an example of the `pwd` command. The example displays the current directory.

Example 12–25 Using pwd

```
ASMCMD [+data/orcl/datafile] > pwd
+data/orcl/datafile
```

rm

Purpose

Deletes the specified Oracle ASM files and directories.

Syntax and Description

```
rm [-f|-r] name [name...]
```

Table 12–24 Options for the rm command

Option	Description
-r	Recursively deletes files and subdirectories.
-f	Forces the deletion of files and subdirectories.
name	Name of the file, pattern, or directory you want to remove.

If *name* is a file or alias, then the `rm` command can delete the file or alias only if it is not currently in use. If *name* is a directory, then the `rm` command can delete it only if it is empty (unless the `-r` flag is used) and it is not a system-generated directory. If *name*

is an alias, then the `rm` command deletes both the alias and the file to which the alias refers. To delete only an alias and retain the file that the alias references, use the `rmalias` command.

Note: When you delete all of the files in a system-created directory, the directory is removed. If the parent directories are empty, all of the parent directories are also removed.

name can contain wildcard characters. See "[Wildcard Characters](#)" on page 12-5.

If you use a wildcard, the `rm` command deletes all of the matches except nonempty directories, unless you use the `-r` flag. To recursively delete, use the `-r` flag. With `-r` option you can delete a nonempty directory, including all files and directories in it and in the entire directory tree underneath it. If you use the `-r` flag or a wildcard character, then the `rm` command prompts you to confirm the deletion before proceeding, unless you specify the `-f` flag.

If a wildcard character matches an alias or a system-generated file that has an alias, then both the alias and the system-generated file that it references are deleted. When using the `-r` flag, either the system-generated file or the alias must be present in the directory in which you run the `rm` command.

For example, if you have a user alias, `+data/dir1/file.alias` that points to `+data/orcl/DATAFILE/System.256.146589651`, then running the `rm -r +data/dir1` command removes the `+data/dir1/file.alias` and `+data/orcl/DATAFILE/System.256.146589651`.

Example

The following are examples of the `rm` command. The first example deletes the `myexamples.bak` file. The second example removes the `subdir2` directory and its contents.

Example 12-26 Using rm

```
ASMCMD [+data/orcl/datafile] > rm myexamples.bak

ASMCMD [+data] > rm -r subdir2
You may delete multiple files and/or directories.
Are you sure? (y/n) y
```

rmalias

Purpose

Removes the specified aliases, retaining the files that the aliases reference.

Syntax and Description

```
rmalias [-r] alias [alias..]
```

Table 12-25 Options for the *rmalias* command

Option	Description
<code>-r</code>	Recursively removes aliases.
<i>alias</i>	Alias for the file name or directory.

The `-r` flag enables you to remove all of the aliases in the current directory and in the entire directory tree beneath the current directory. If any user-created directories become empty as a result of deleting aliases, they are also deleted. Files and directories created by the system are not deleted.

Example

The following is an example of the `rmalias` command. The example deletes the alias `sysaux.f`, retaining the data file that it references.

Example 12–27 Using `rmalias`

```
ASMCMD [+data/orcl/datafile] > rmalias sysaux.f
```

ASMCMD Disk Group Management Commands

This section describes the ASMCMD disk group management commands.

[Table 12–26](#) provides a summary of the disk group management commands.

Table 12–26 Summary of ASMCMD Disk Group Management Commands

Command	Description
<code>chdg</code>	Changes a disk group (add, drop, or rebalance).
<code>chkdg</code>	Checks or repairs a disk group.
<code>dropdg</code>	Drops a disk group.
<code>iostat</code>	Displays I/O statistics for disks.
<code>lsattr</code>	Lists the attribute of a disk group.
<code>lsdg</code>	Lists disk groups and their information.
<code>lsdsk</code>	Lists disks Oracle ASM disks.
<code>lsod</code>	Lists open devices.
<code>md_backup</code>	Creates a backup of the metadata of mounted disk groups.
<code>md_restore</code>	Restores disk groups from a backup of the metadata.
<code>mkdg</code>	Creates a disk group.
<code>mount</code>	Mounts a disk group.
<code>offline</code>	Offline a disk or a failure group.
<code>online</code>	Onlines a disk or a failure group.
<code>rebal</code>	Rebalances a disk group.
<code>remap</code>	Relocates data in a range of physical blocks on a disk.
<code>setattr</code>	Sets attributes in a disk group.
<code>umount</code>	Dismounts a disk group.

chdg

Purpose

Changes a disk group (adds disks, drops disks, or rebalances) based on an XML configuration file.

Syntax and Description

```
chdg { config_file.xml | 'contents_of_xml_file' }
```

Table 12–27 Options for the chdg Command

Option	Description
<i>config_file</i>	Name of the XML file that contains the changes for the disk group. chdg searches for the XML file in the directory where ASMCMD was started unless a path is specified. For examples of the valid tags and XML configuration file, see Example 12–28 and Example 12–29 .
<i>contents_of_xml_file</i>	The XML script enclosed in single quotations.

chdg modifies a disk group based on an XML configuration file. The modification includes adding or deleting disks from an existing disk group, and the setting rebalance power level. The power level can set from 0 to the maximum of 11, the same values as the ASM_POWER_LIMIT initialization parameter. For information about the initialization parameter, see "[ASM_POWER_LIMIT](#)" on page 3-8.

When adding disks to a disk group, the diskstring must be specified similar to the ASM_DISKSTRING initialization parameter. For information about the initialization parameter, see "[ASM_DISKSTRING](#)" on page 3-7.

The failure groups are optional parameters. The default causes every disk to belong to a its own failure group. For information about failure groups, see "[Oracle ASM Failure Groups](#)" on page 4-24.

Dropping disks from a disk group can be performed through this operation. An individual disk can be referenced by its Oracle ASM disk name. A the set of disks that belong to a failure group can be specified by the failure group name. For information about dropping disks, see "[Dropping Disks from Disk Groups](#)" on page 4-14.

You can resize a disk inside a disk group with chdg. The resize operation fails if there is not enough space for storing data after the resize. For information about resizing disks, see "[Resizing Disks in Disk Groups](#)" on page 4-16.

[Example 12–28](#) shows the basic structure and the valid tags with their respective attributes for the chdg XML configuration file.

Example 12–28 Tags for the chdg XML Configuration Template

```
<chdg> update disk clause (add/delete disks/failure groups)
      name      disk group to change
      power     power to perform rebalance

<add>  items to add are placed here
</add>

<drop> items to drop are placed here
</drop>

<fg>  failure group
      name      failure group name
</fg>

<dsk> disk
      name      disk name
      path      disk path
```

```

        size          size of the disk to add
</dsk>

</chdg>

```

For information about creating a disk group with ASMCMD `mkdgm`, see "[mkdgm](#)" on page 12-40. For information about altering disk groups, see "[Altering Disk Groups](#)" on page 4-11. For information about XML, see <http://www.oracle.com/technology/tech/xml/index.html>

Example

The following is an example of an XML configuration file for `chdg`. This XML file alters the disk group named `data`. The failure group `fg1` is dropped and the disk `data_0001` is also dropped. The `/dev/disk5` disk is added to failure group `fg2`. The rebalance power level is set to 3.

Example 12-29 `chdg` sample XML configuration file

```

<chdg name="data" power="3">
  <drop>
    <fg name="fg1"></fg>
    <dsk name="data_0001"/>
  </drop>
  <add>
    <fg name="fg2">
      <dsk string="/dev/disk5"/>
    </fg>
  </add>
</chdg>

```

The following are examples of the `chdg` command with the configuration file or configuration information on the command line.

Example 12-30 Using `chdg`

```

ASMCMD [+] > chdg data_config.xml

ASMCMD [+] > chdg '<chdg name="data" power="3">
  <drop><fg name="fg1"></fg><dsk name="data_0001"/></drop>
  <add><fg name="fg2"><dsk string="/dev/disk5"/></fg></add></chdg>'

```

chkdgm

Purpose

Checks or repairs the metadata of a disk group.

Syntax and Description

```
chkdgm [--repair] diskgroup
```

Table 12-28 Options for the `chkdgm` Command

Option	Description
<code>--repair</code>	Repairs the disk group.
<i>diskgroup</i>	Name of disk group to check or repair.

chkdsk checks the metadata of a disk group for errors and optionally repairs the errors.

Example

The following is an example of the chkdsk command used to check and repair the data disk group.

Example 12-31 Using chkdsk

```
ASMCMD [+] > chkdsk --repair data
```

dropdg

Purpose

Drops a disk group.

Syntax and Description

```
dropdg [-r -f] [-r] diskgroup
```

Table 12-29 Options for the dropdg Command

Option	Description
-f	Force the operation. Only applicable if the disk group cannot be mounted.
-r	Recursive, include contents.
<i>diskgroup</i>	Name of disk group to drop.

dropdg drops an existing disk group. The disk group cannot be mounted on more than one node.

Example

These are examples of the use of dropdg. The first example forces the drop of the disk group data, including any data in the disk group. The second example drops the disk group fra, including any data in the disk group.

Example 12-32 Using dropdg

```
ASMCMD [+] > dropdg -r -f data
```

```
ASMCMD [+] > dropdg -r fra
```

iostat

Purpose

Displays I/O statistics for Oracle ASM disks in mounted disk groups.

Syntax and Description

```
iostat [-eth][--io][--region] [-G diskgroup] [interval]
```

iostat lists disk group statistics using the V\$ASM_DISK_IOSTAT view.

Table 12–30 Options for the iostat command

Option	Description
-e	Displays error statistics (Read_Err, Write_Err).
-G <i>diskgroup</i>	Displays statistics for the disk group name.
-H	Suppresses column headings.
--io	Displays information in number of I/Os, instead of bytes.
-t	Displays time statistics (Read_Time, Write_Time).
--region	Displays information for cold and hot disk regions (Cold_Reads, Cold_Writes, Hot_Reads, Hot_Writes).
<i>interval</i>	Refreshes the statistics display based on the interval value (seconds).

Table 12–31 shows the statistics for a disk group. To view the complete set of statistics for a disk group, use the V\$ASM_DISK_IOSTAT view.

Table 12–31 Attribute descriptions for iostat command output

Attribute Name	Description
Group_Name	Name of the disk group.
Dsk_Name	Name of the disk.
Reads	Total number of bytes read from the disk. If the --io option is entered, then the value is displayed as number of I/Os.
Writes	Total number of bytes written from the disk. If the --io option is entered, then the value is displayed as number of I/Os.
Cold_Reads	Total number of bytes read from the cold disk region. If the --io option is entered, then the value is displayed as number of I/Os.
Cold_Writes	Total number of bytes written from the cold disk region. If the --io option is entered, then the value is displayed as number of I/Os.
Hot_Reads	Total number of bytes read from the hot disk region. If the --io option is entered, then the value is displayed as number of I/Os.
Hot_Writes	Total number of bytes written from the hot disk region. If the --io option is entered, then the value is displayed as number of I/Os.
Read_Err	Total number of failed I/O read requests for the disk.
Write_Err	Total number of failed I/O write requests for the disk.
Read_Time	Total I/O time (in hundredths of a second) for read requests for the disk if the TIMED_STATISTICS initialization parameter is set to TRUE (0 if set to FALSE).
Write_Time	Total I/O time (in hundredths of a second) for write requests for the disk if the TIMED_STATISTICS initialization parameter is set to TRUE (0 if set to FALSE).

Example

The following are examples of the `iostat` command. The first example displays disk I/O statistics for the `data` disk group in total number of bytes. The second example displays disk I/O statistics for the `data` disk group in total number of I/O operations.

Example 12–33 Using iostat

```
ASMCMD [+] > iostat -G data
```



```

Group_Name Dsk_Name Reads Writes
DATA DATA_0000 180488192 473707520
DATA DATA_0001 1089585152 469538816
DATA DATA_0002 191648256 489570304
DATA DATA_0003 175724032 424845824
DATA DATA_0004 183421952 781429248
DATA DATA_0005 1102540800 855269888
DATA DATA_0006 171290624 447662592
DATA DATA_0007 172281856 361337344
DATA DATA_0008 173225472 390840320
DATA DATA_0009 288497152 838680576
DATA DATA_0010 196657152 375764480
DATA DATA_0011 436420096 356003840

```

```

ASMCMD [+] > iostat --io -G data
Group_Name Dsk_Name Reads Writes
DATA DATA_0000 2801 34918
DATA DATA_0001 58301 35700
DATA DATA_0002 3320 36345
DATA DATA_0003 2816 10629
DATA DATA_0004 2883 34850
DATA DATA_0005 59306 38097
DATA DATA_0006 2151 10129
DATA DATA_0007 2686 10376
DATA DATA_0008 2105 8955
DATA DATA_0009 9121 36713
DATA DATA_0010 3557 8596
DATA DATA_0011 17458 9269

```

lsattr

Purpose

Lists the attributes of a disk group.

Syntax and Description

```
lsattr [-G diskgroup] [-Hlm] [pattern]
```

Table 12-32 Options for the lsattr command

Option	Description
-G <i>diskgroup</i>	Disk group name.
-H	Suppresses column headings.
-l	Display names with values.
-m	Displays additional information, such as the RO and Sys columns.
<i>pattern</i>	Display the attributes that contain pattern expression.

Information about disk group attributes is retrieved from the V\$ASM_ATTRIBUTE view. For information about disk group attributes, see ["Disk Group Attributes"](#) on page 4-2.

The RO (read-only) column identifies those attributes that can only be set when a disk group is created. The Sys column identifies those attributes that are system-created.

To display information about the disk group template attributes, see "lstmpl" on page 12-48.

Example

The following are examples of the `lsattr` command. The first displays information about all attributes for the `data` disk group. The second example displays only those attributes with names containing the string `compat` for the `fra` disk group.

Example 12-34 Using lsattr

```
ASMCMD [+] > lsattr -l -G data
Name                               Value
access_control.enabled             FALSE
access_control.umask               066
au_size                            4194304
cell.smart_scan_capable            FALSE
compatible.advm                    11.2.0.0.0
compatible.asm                     11.2.0.0.0
compatible.rdbms                   10.1.0.0.0
disk_repair_time                   3.6h
sector_size                         512

ASMCMD [+] > lsattr -l -G fra %compat*
Name                               Value
compatible.asm                     11.2.0.0.0
compatible.rdbms                   10.1.0.0.0
```

lsdg

Purpose

Lists disk groups and their information. `lsdg` queries `V$ASM_DISKGROUP_STAT` by default. If the `--discovery` flag is specified, the `V$ASM_DISKGROUP` is queried instead. The output also includes notification of any current rebalance operation for a disk group. If a disk group is specified, then `lsdg` returns only information about that disk group.

Syntax and Description

```
lsdg [-gH] [--discovery] [pattern]
```

Table 12-33 lists the options for the `lsdg` command.

Table 12-33 Options for the lsdg command

Option	Description
(none)	Displays the disk group attributes listed in Table 12-34.
<code>--discovery</code>	Selects from <code>V\$ASM_DISKGROUP</code> , or from <code>GV\$ASM_DISKGROUP</code> if the <code>-g</code> flag is also specified. This option is always enabled if the Oracle ASM instance is version 10.1 or earlier. This flag is disregarded if <code>lsdsk</code> is running in non-connected mode.
<code>-g</code>	Selects from <code>GV\$ASM_DISKGROUP_STAT</code> , or from <code>GV\$ASM_DISKGROUP</code> if the <code>--discovery</code> flag is also specified. <code>GV\$ASM_DISKGROUP.INST_ID</code> is included in the output. The <code>REBAL</code> column of the <code>GV\$ASM_OPERATION</code> view is also included in the output.
<code>-H</code>	Suppresses column headings.

Table 12–33 (Cont.) Options for the lsdg command

Option	Description
<i>pattern</i>	Returns only information about the specified disk group or disk groups that match the supplied pattern. See "Wildcard Characters" on page 12-5.

Table 12–34 shows the attributes for each disk group. To view the complete set of attributes for a disk group, use the V\$ASM_DISKGROUP_STAT or V\$ASM_DISKGROUP view.

See Also: *Oracle Database Reference* for descriptions of disk group information displayed in the V\$ASM_DISKGROUP view

Table 12–34 Attribute descriptions for lsdg command output

Attribute Name	Description
State	State of the disk group. Values include BROKEN, CONNECTED, DISMOUNTED, MOUNTED, QUIESCING, and UNKNOWN.
Type	Disk group redundancy (NORMAL, HIGH, EXTERNAL).
Rebal	Y if a rebalance operation is in progress.
Sector	Sector size in bytes.
Block	Block size in bytes.
AU	Allocation unit size in bytes.
Total_MB	Size of the disk group in megabytes.
Free_MB	Free space in the disk group in megabytes, without regard to redundancy. From the V\$ASM_DISKGROUP view.
Req_mir_free_MB	Amount of space that must be available in the disk group to restore full redundancy after the most severe failure that can be tolerated by the disk group. This is the REQUIRED_MIRROR_FREE_MB column from the V\$ASM_DISKGROUP view.
Usable_file_MB	Amount of free space, adjusted for mirroring, that is available for new files. From the V\$ASM_DISKGROUP view.
Offline_disks	Number of offline disks in the disk group. Offline disks are eventually dropped.
Voting_files	Specifies whether the disk group contains voting disks (Y or N).
Name	Disk group name.

Example

The following example lists the attributes of the data disk group.

Example 12–35 Using lsdg

```
ASMCMD [+] > lsdg data
State   Type   Rebal  Sector  Block      AU   Total_MB  Free_MB  Req_mir_free_MB  Usable_file_MB
MOUNTED NORMAL N        512    4096   4194304    12288    8835           1117           3859

(continued)
Offline_disks  Voting_files  Name
              0              N  DATA
```

lsdisk

Purpose

List Oracle ASM disks.

Syntax and Description

```
lsdisk [-kptgMHI] [-G diskgroup] [--member|--candidate]
      [--discovery][--statistics][pattern]
```

Table 12–35 lists the options for the `lsdisk` command.

Table 12–35 Options for the `lsdisk` command

Option	Description
(none)	Displays the <code>PATH</code> column of the <code>V\$ASM_DISK</code> view.
-k	Displays the <code>TOTAL_MB</code> , <code>FREE_MB</code> , <code>OS_MB</code> , <code>NAME</code> , <code>FAILGROUP</code> , <code>LIBRARY</code> , <code>LABEL</code> , <code>UDID</code> , <code>PRODUCT</code> , <code>REDUNDANCY</code> , and <code>PATH</code> columns of the <code>V\$ASM_DISK</code> view.
--statistics	Displays the <code>READS</code> , <code>WRITES</code> , <code>READ_ERRS</code> , <code>WRITE_ERRS</code> , <code>READ_TIME</code> , <code>WRITE_TIME</code> , <code>BYTES_READ</code> , <code>BYTES_WRITTEN</code> , and the <code>PATH</code> columns of the <code>V\$ASM_DISK</code> view.
-p	Displays the <code>GROUP_NUMBER</code> , <code>DISK_NUMBER</code> , <code>INCARNATION</code> , <code>MOUNT_STATUS</code> , <code>HEADER_STATUS</code> , <code>MODE_STATUS</code> , <code>STATE</code> , and the <code>PATH</code> columns of the <code>V\$ASM_DISK</code> view.
-t	Displays the <code>CREATE_DATE</code> , <code>MOUNT_DATE</code> , <code>REPAIR_TIMER</code> , and the <code>PATH</code> columns of the <code>V\$ASM_DISK</code> view.
-g	Selects from <code>GV\$ASM_DISK_STAT</code> , or from <code>GV\$ASM_DISK</code> if the <code>--discovery</code> flag is also specified. <code>GV\$ASM_DISK.INST_ID</code> is included in the output.
--discovery	Selects from <code>V\$ASM_DISK</code> , or from <code>GV\$ASM_DISK</code> if the <code>-g</code> flag is also specified. This option is always enabled if the Oracle ASM instance is version 10.1 or earlier. This flag is disregarded if <code>lsdisk</code> is running in non-connected mode.
-H	Suppresses column headings.
-I	Scans disk headers for information rather than extracting the information from an Oracle ASM instance. This option forces non-connected mode.
-G	Restricts results to only those disks that belong to the group specified by <i>diskgroup</i> .
-M	Displays the disks that are visible to some but not all active instances. These are disks that, if included in a disk group, cause the mount of that disk group to fail on the instances where the disks are not visible.
--candidate	Restricts results to only disks having membership status equal to <code>CANDIDATE</code> .
--member	Restricts results to only disks having membership status equal to <code>MEMBER</code> .
<i>pattern</i>	Returns only information about the specified disks that match the supplied pattern.

The `lsdisk` command can run in connected or non-connected mode. The connected mode is always attempted first. The `-I` option forces non-connected mode.

- In connected mode, `lsdisk` uses the `V$ASM_DISK_STAT` and `V$ASM_DISK` dynamic views to retrieve disk information. The `V$ASM_DISK_STAT` view is used by default.
- In non-connected mode, `lsdisk` scans disk headers to retrieve disk information. Some information is not available in this mode and some options are not valid combinations with this mode.

Note: The non-connected mode is not supported on Windows.

pattern restricts the output to only disks that match the pattern specified. Wild-card characters and slashes (/ or \) can be part of the pattern. *pattern* should be specified as the last option for the command. For information about wildcards, see "[Wildcard Characters](#)" on page 12-5.

The `-k`, `-p`, `-t`, and `--statistics` options modify how much information is displayed for each disk. If any combination of the options are specified, then the output shows the union of the attributes associated with each flag.

Example

The following are examples of the `lsdisk` command. The first and second examples list information about disks in the data disk group. The third example lists information about candidate disks.

Example 12–36 Using `lsdisk`

```
ASMCMD [+] > lsdisk -t -G data
```

Create_Date	Mount_Date	Repair_Timer	Path
13-JUL-09	13-JUL-09	0	/devices/diska1
13-JUL-09	13-JUL-09	0	/devices/diska2
13-JUL-09	13-JUL-09	0	/devices/diska3
13-JUL-09	13-JUL-09	0	/devices/diskb1
13-JUL-09	13-JUL-09	0	/devices/diskb2
13-JUL-09	13-JUL-09	0	/devices/diskb3
13-JUL-09	13-JUL-09	0	/devices/diskc1
13-JUL-09	13-JUL-09	0	/devices/diskc2
...			

```
ASMCMD [+] > lsdisk -p -G data /devices/diska*
```

Group_Num	Disk_Num	Incarn	Mount_Stat	Header_Stat	Mode_Stat	State	Path
1	0	2105454210	CACHED	MEMBER	ONLINE	NORMAL	/devices/diska1
1	1	2105454199	CACHED	MEMBER	ONLINE	NORMAL	/devices/diska2
1	2	2105454205	CACHED	MEMBER	ONLINE	NORMAL	/devices/diska3

```
ASMCMD [+] > lsdisk --candidate -p
```

Group_Num	Disk_Num	Incarn	Mount_Stat	Header_Stat	Mode_Stat	State	Path
0	5	2105454171	CLOSED	CANDIDATE	ONLINE	NORMAL	/devices/diske1
0	25	2105454191	CLOSED	CANDIDATE	ONLINE	NORMAL	/devices/diske2
0	18	2105454184	CLOSED	CANDIDATE	ONLINE	NORMAL	/devices/diske3
0	31	2105454197	CLOSED	CANDIDATE	ONLINE	NORMAL	/devices/diskk1
0	21	2105454187	CLOSED	CANDIDATE	ONLINE	NORMAL	/devices/diskk2
0	26	2105454192	CLOSED	CANDIDATE	ONLINE	NORMAL	/devices/diskk3
0	14	2105454180	CLOSED	CANDIDATE	ONLINE	NORMAL	/devices/diskl1
...							

lsod

Purpose

Lists the open ASM disks.

Syntax and Description

```
lsod [-H] [-G diskgroup] [--process process] [pattern]
```

Table 12–36 Options for the lsod command

Option	Description
-H	Suppresses column header information from the output.
-G <i>diskgroup</i>	Specifies the disk group that contains the open disks.
--process <i>process</i>	Specifies a pattern to filter the list of processes.
<i>pattern</i>	Specifies a pattern to filter the list of disks.

The rebalance operation (RBAL) opens a disk both globally and locally so the same disk may be listed twice in the output for the RBAL process.

Example

The following are examples of the `lsod` command. The first example lists the open devices associated with the `data` disk group and the `LGWR` process. The second example lists the open devices associated with the `LGWR` process for disks that match the `diska` pattern.

Example 12–37 Using lsod

```
ASMCMD [+] > lsod -G data --process LGWR
Instance Process                OSPID Path
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diska1
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diska2
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diska3
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diskb1
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diskb2
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diskb3
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diskd1

ASMCMD [+] > lsod --process LGWR diska
Instance Process                OSPID Path
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diska1
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diska2
1         oracle@dadvmn0652 (LGWR) 26593 /devices/diska3
```

For another example of the `lsod` command, see [Example 12–2](#) on page 12-8.

md_backup

Purpose

The `md_backup` command creates a backup file containing metadata for one or more disk groups.

Volume and Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file system information is not backed up.

Syntax and Description

```
md_backup backup_file
      [-G 'diskgroup[,diskgroup,...]']
```

Table 12–37 describes the options for the md_backup command.

Table 12–37 Options for the md_backup Command

Option	Description
<i>backup_file</i>	Specifies the backup file in which you want to store the metadata.
<i>-G diskgroup</i>	Specifies the disk group name of the disk group that must be backed up

By default all the mounted disk groups are included in the backup file, which is saved in the current working directory.

Example

The first example shows the use of the backup command when you run it without options. This example backs up all of the mounted disk groups and creates the backup image in the current working directory. The second example creates a backup of disk group *data* and *fra*. The backup that this example creates is saved in the */tmp/dgbackup20090714* file.

Example 12–38 Using md_backup

```
ASMCMD [+] > md_backup /tmp/dgbackup20090716

ASMCMD [+] > md_backup /tmp/dgbackup20090716 -G data, fra
Disk group metadata to be backed up: DATA
Disk group metadata to be backed up: FRA
Current alias directory path: ASM/ASMPARAMETERFILE
Current alias directory path: ORCL/DATAFILE
Current alias directory path: ORCL/TEMPFILE
Current alias directory path: ORCL/CONTROLFILE
Current alias directory path: ORCL/PARAMETERFILE
Current alias directory path: ASM
Current alias directory path: ORCL/ONLINELOG
Current alias directory path: ORCL
Current alias directory path: ORCL/CONTROLFILE
Current alias directory path: ORCL/ARCHIVELOG/2009_07_13
Current alias directory path: ORCL/BACKUPSET/2009_07_14
Current alias directory path: ORCL/ARCHIVELOG/2009_07_14
Current alias directory path: ORCL
Current alias directory path: ORCL/DATAFILE
Current alias directory path: ORCL/ARCHIVELOG
Current alias directory path: ORCL/BACKUPSET
Current alias directory path: ORCL/ONLINELOG
```

md_restore

Purpose

This command restores a disk group metadata backup.

Syntax and Description

```
md_restore backup_file [--silent]
    [--full|--nodg|--newdg -o 'old_diskgroup:new_diskgroup [, ...]']
    [-S sql_script_file] [-G 'diskgroup[, diskgroup..]']
```

Table 12–38 describes the options for the md_restore command.

Table 12–38 Options for the md_restore command

Option	Description
<i>backup_file</i>	Reads the metadata information from <i>backup_file</i> .
--silent	Ignore errors. Typically, if md_restore encounters an error, it stops. Specifying this flag ignores any errors.
--full	Specifies to create a disk group and restore metadata.
--nodg	Specifies to restore metadata only.
--newdg -o <i>old_diskgroup:new_diskgroup</i>]	Specifies to create a disk group with a different name when restoring metadata. The -o option is required with --newdg.
-S <i>sql_script_file</i>	Write SQL commands to the specified SQL script file instead of executing the commands.
-G <i>diskgroup</i>	Select the disk groups to be restored. If no disk groups are defined, then all disk groups are restored.

Example

The first example restores the disk group data from the backup script and creates a copy. The second example takes an existing disk group data and restores its metadata. The third example restores disk group data completely but the new disk group that is created is called data2. The fourth example restores from the backup file after applying the overrides defined in the override.sql script file.

Example 12–39 Using md_restore

```
ASMCMD [+] > md_restore --full -G data --silent /tmp/dgbackup20090714
ASMCMD [+] > md_restore --nodg -G data --silent /tmp/dgbackup20090714
ASMCMD [+] > md_restore --newdg -o 'data:data2' --silent /tmp/dgbackup20090714
ASMCMD [+] > md_restore -S override.sql --silent /tmp/dgbackup20090714
```

mkdg

Purpose

Creates a disk group based on an XML configuration file.

Syntax and Description

```
mkdg { config_file.xml | 'contents_of_xml_file' }
```


Table 12–39 Options for the mkdg Command

Option	Description
<code>config_file</code>	Name of the XML file that contains the configuration for the new disk group. <code>mkdg</code> searches for the XML file in the directory where ASMCMD was started unless a path is specified. For examples of the valid tags and XML configuration file, see Example 12–40 and Example 12–41 .
<code>contents_of_xml_file</code>	The XML script enclosed in single quotations.

`mkdg` creates a new disk group with an XML configuration file that specifies the name of the disk group, redundancy, attributes, and paths of the disks that form the disk group. Redundancy is an optional parameter; the default is normal redundancy. For some types of redundancy, disks are required to be gathered into failure groups. In the case that failure groups are not specified, every disk will be in its own failure group.

It is possible to set some of the disk group attribute values during disk group creation. Some attributes, such as `AU_SIZE` and `SECTOR_SIZE`, can be set only during disk group creation. For more information about disk groups attributes, "[Disk Group Attributes](#)" on page 4-2.

The default disk group compatibility settings are 10.1 for Oracle ASM compatibility, 10.1 for database compatibility, and no value for Oracle ADVM compatibility. For information about disk group compatibility attributes, see "[Disk Group Compatibility](#)" on page 4-31.

[Example 12–40](#) shows the basic structure and the valid tags with their respective attributes for the `mkdg` XML configuration file.

Example 12–40 Tags for mkdg XML Configuration File

```
<dg> disk group
      name          disk group name
      redundancy    normal, external, high

<fg> failure group
      name          failure group name
</fg>

<disk> disk
      name          disk name
      path          disk path
      size          size of the disk to add
</disk>

<a>   attribute
      name          attribute name
      value         attribute value
</a>

</dg>
```

For information about altering a disk group with ASMCMD `chdg`, see "[chdg](#)" on page 12-28. For information about creating a disk group, see "[Creating Disk Groups](#)" on page 4-3. For information about XML, see <http://www.oracle.com/technology/tech/xml/index.html>

Example

The following is an example of an XML configuration file for `mkdgm`. The configuration file creates a disk group named `data` with `normal` redundancy. Two failure groups, `fg1` and `fg2`, are created, each with two disks identified by associated disk strings. The disk group compatibility attributes are all set to `11.2`.

Example 12–41 `mkdgm` sample XML configuration file

```
<dg name="data" redundancy="normal">
  <fg name="fg1">
    <dsk string="/dev/disk1"/>
    <dsk string="/dev/disk2"/>
  </fg>
  <fg name="fg2">
    <dsk string="/dev/disk3"/>
    <dsk string="/dev/disk4"/>
  </fg>
  <a name="compatible.asm" value="11.2"/>
  <a name="compatible.rdbms" value="11.2"/>
  <a name="compatible.advm" value="11.2"/>
</dg>
```

The following are examples of the `mkdgm` command. The first example executes `mkdgm` with an XML configuration file in the directory where ASMCMD was started. The second examples executes `mkdgm` using information on the command line.

Example 12–42 Using `mkdgm`

```
ASMCMD [+] > mkdgm data_config.xml

ASMCMD [+] > mkdgm '<dg name="data"><dsk path="/dev/disk*"/></dg>'
```

mount

Purpose

Mounts a disk group.

Syntax and Description

```
mount [--restrict] { [-a] | [-f] diskgroup [ diskgroup ... ] }
```

Table 12–40 Options for the `mount` Command

Option	Description
<i>diskgroup</i>	Name of the disk group.
-a	Mounts all disk groups.
--restrict	Mounts in restricted mode.
-f	Forces the mount operation.

This operation mounts one or more disk groups. A disk group can be mounted with or without force or restricted options. For more information about mounting disk groups, see "[Mounting and Dismounting Disk Groups](#)" on page 4-42.

Example

The following are examples of the `mount` command showing the use of the `force`, `restrict`, and all options.

Example 12-43 Using mount

```
ASMCMD [+] > mount -f data

ASMCMD [+] > mount --restrict data

ASMCMD [+] > mount -a
```

offline

Purpose

Offline disks or failure groups that belong to a disk group.

Syntax and Description

```
offline -G diskgroup
        { -F failgroup | -D disk }
        [-t {minutes | hours}]
```

Table 12-41 Options for the offline command

Option	Description
-G <i>diskgroup</i>	Disk group name.
-F <i>failgroup</i>	Failure group name.
-D <i>disk</i>	Specifies a single disk name.
-t <i>minutes</i> <i>hours</i>	Specifies the time before the specified disk is dropped as <i>mm</i> or <i>nh</i> , where <i>m</i> specifies minutes and <i>h</i> specifies hours. The default unit is hours.

When a failure group is specified, this implies all the disks that belong to it should be offlined.

Example

The following are examples of the `offline` command. The first example offlines the `failgroup1` failure group of the `data` disk group. The second example offlines the `data_0001` disk of the `data` disk group with a time of 1.5 hours before the disk is dropped.

Example 12-44 Using offline

```
ASMCMD [+] > offline -G data -F failgroup1

ASMCMD [+] > offline -G data -D data_0001 -t 1.5h
```

online

Purpose

Online all disks, a single disk, or a failure group that belongs to a disk group.

Syntax and Description

`online { [-a] -G diskgroup | -F failgroup | -D disk } [-w]`

Table 12–42 Options for the online command

Option	Description
-a	Online all offline disks in the disk group.
-G <i>diskgroup</i>	Disk group name.
-F <i>failgroup</i>	Failure group name.
-D <i>disk</i>	Disk name.
-w	Wait option. Causes ASMCMD to wait for the disk group to be rebalanced before returning control to the user. The default is not waiting.

When a failure group is specified, this implies all the disks that belong to it should be online.

Example

The following are examples of the `online` command. The first example online all disks in the `failgroup1` failure group of the `data` disk group with the `wait` option enabled. The second example online the `data_0001` disk in the `data` disk group.

Example 12–45 Using online

```
ASMCMD [+] > online -G data -F failgroup1 -w
ASMCMD [+] > online -G data -D data_0001
```

rebal

Purpose

Rebalances a disk group.

Syntax and Description

`rebal [--power power] [-w] diskgroup`

Table 12–43 Options for the rebal command

Option	Description
<i>diskgroup</i>	Disk group name.
--power <i>power</i>	Power setting (0 to 11).
-w	Wait option. Causes ASMCMD to wait for the disk group to be rebalanced before returning control to the user. The default is not waiting.

The power level can be set from 0 to 11. A value of 0 disables rebalancing. If the rebalance power is not specified, the value defaults to the setting of the `ASM_POWER_LIMIT` initialization parameter. For information about the power level, see "[ASM_POWER_LIMIT](#)" on page 3-8 and "[Tuning Rebalance Operations](#)" on page 4-18.

You can determine if a rebalance operation is occurring with the ASMCMD `lsop` command. See "lsop" on page 12-12.

For more information about rebalancing a disk group, see "Manually Rebalancing Disk Groups" on page 4-17.

Example

The following is an example of the `rebal` command that rebalances the `fra` disk group with a power level set to 4.

Example 12-46 Using rebal

```
ASMCMD [+] > rebal --power 4 fra

ASMCMD [+] > lsop
Group_Name Dsk_Num State Power
FRA        REBAL  RUN    4
```

remap

Purpose

Marks a range of blocks as unusable on the disk and relocates any data allocated in that range.

Syntax and Description

```
remap diskgroup disk block_range
```

Table 12-44 Options for the remap command

Option	Description
<i>diskgroup</i>	Disk group name in which a disk must have data relocated.
<i>disk</i>	Name of the disk that must have data relocated. The name must match the NAME column in the V\$ASM_DISK view.
<i>block_range</i>	Range of physical blocks to relocate in the format <i>start_range_number-end_range_number</i> .

The `remap` command only relocates blocks. It does not correct or repair blocks that contain corrupted contents.

The command uses a physical block size based on the `SECTOR_SIZE` disk group attribute.

Examples

The first example remaps blocks 5000 through 5999 for disk `DATA_0001` in disk group `DATA`. The second example remaps blocks 6230 through 6339 for disk `FRA_0002` in disk group `FRA`.

Example 12-47 Using remap

```
ASMCMD [+] > remap DATA DATA_0001 5000-5999

ASMCMD [+] > remap FRA FRA_0002 6230-6339
```

setattr

Purpose

Sets the attributes for an Oracle ASM disk group.

Syntax and Description

```
setattr -G diskgroup attribute_name attribute_value
```

Table 12–45 Options for the setattr command

Option	Description
<i>-G diskgroup</i>	Disk group name.
<i>attribute_name</i>	Name of the attribute.
<i>attribute_value</i>	Value of the attribute.

For information about disk group attributes, see "[Disk Group Attributes](#)" on page 4-2.

Example

The following are examples of the `setattr` command. The first example sets the disk group attribute `COMPATIBLE.ASM` to `11.2` for the `data` disk group. The second example sets the disk group attribute `COMPATIBLE.RDBMS` to `11.2` for the `data` disk group.

Example 12–48 Using setattr

```
ASMCMD [+] > setattr -G data compatible.asm 11.2.0.0.0
```

```
ASMCMD [+] > setattr -G data compatible.rdbms 11.2.0.0.0
```

umount

Purpose

Dismounts a disk group.

Syntax and Description

```
umount { -a | [-f] diskgroup }
```

Table 12–46 Options for the umount Command

Option	Description
<i>diskgroup</i>	Name of the disk group.
<i>-a</i>	Dismounts all mounted disk groups. These disk groups are listed in the output of the <code>V\$ASM_DISKGROUP</code> view.
<i>-f</i>	Forces the dismount operation.

Example

The following are examples of the `umount` command. The first example dismounts all disk groups mounted on the Oracle ASM instance. The second example forces the dismount of the `data` disk group.

Example 12–49 Using umount

```
ASMCMD [+] > umount -a
```

```
ASMCMD [+] > umount -f data
```

ASMCMD Template Management Commands

This section describes the ASMCMD template management commands.

[Table 12–47](#) provides a summary of the template management commands.

Table 12–47 Summary of ASMCMD Template Management Commands

Command	Description
chtml	Changes the attributes of a template.
lstmpl	Lists all templates or the templates for a specified disk group.
mktmpl	Adds a template to a disk group.
rmtmpl	Removes a disk group template.

chtml

Purpose

Changes the attributes of a template.

Syntax and Description

```
chtml -G diskgroup { [--striping {coarse|fine}]
    [--redundancy {high|mirror|unprotected} ] [--primary {hot|cold}]
    [--secondary {hot|cold}]} template
```

Table 12–48 Options for the chhtml Command

Option	Description
<i>-G diskgroup</i>	Name of the disk group.
<i>template</i>	Name of the template to change.
--striping {coarse fine}	Striping specification, either coarse or fine).
--redundancy { high mirror unprotected}	Redundancy specification, either high, mirror, or unprotected.
--primary { hot cold }	Intelligent Data Placement specification for primary extents, either hot or cold region.
--secondary { hot cold }	Intelligent Data Placement specification for secondary extents, either hot or cold region.

At least one of these options are required: --striping, --redundancy, --primary, or --secondary.

For information about updating templates, see ["Modifying a Disk Group Template"](#) on page 7-18.

Example

The following is an example of the `chtmpl` command that updates the `mytemplate` template of the `data` disk group. The redundancy attribute is set to `high` and the striping attribute is set to `fine`.

Example 12–50 Using `chtmpl`

```
ASMCMD [+] > chtmpl -G data --redundancy high --striping fine mytemplate
```

lstmpl

Purpose

Lists all templates or the templates for a specified disk group.

Syntax and Description

```
lstmpl [-H1] [-G diskgroup] [pattern]
```

Table 12–49 Options for the `lstmpl` command

Option	Description
<code>-G <i>diskgroup</i></code>	Specifies disk group name.
<code>-H</code>	Suppresses column headings.
<code>-1</code>	Displays all details.
<i>pattern</i>	Displays the templates that match pattern expression.

For information about template attributes, see "[Template Attributes](#)" on page 7-15.

Example

The following is an example of the `lstmpl` command. The example lists all details of the templates in the `data` disk group.

Example 12–51 Using `lstmpl`

```
ASMCMD [+] > lstmpl -1 -G data
```

Group_Name	Group_Num	Name	Stripe	Sys	Redund	PriReg	MirrReg
DATA	1	ARCHIVELOG	COARSE	Y	MIRROR	COLD	COLD
DATA	1	ASMPARAMETERFILE	COARSE	Y	MIRROR	COLD	COLD
DATA	1	AUTOBACKUP	COARSE	Y	MIRROR	COLD	COLD
DATA	1	BACKUPSET	COARSE	Y	MIRROR	COLD	COLD
DATA	1	CHANGETRACKING	COARSE	Y	MIRROR	COLD	COLD
DATA	1	CONTROLFILE	FINE	Y	HIGH	COLD	COLD
DATA	1	DATAFILE	COARSE	Y	MIRROR	COLD	COLD
DATA	1	DATAGUARDCONFIG	COARSE	Y	MIRROR	COLD	COLD
DATA	1	DUMPSET	COARSE	Y	MIRROR	COLD	COLD
DATA	1	FLASHBACK	COARSE	Y	MIRROR	COLD	COLD
DATA	1	MYTEMPLATE	FINE	N	HIGH	COLD	COLD
DATA	1	OCRFILE	COARSE	Y	MIRROR	COLD	COLD
DATA	1	ONLINELOG	COARSE	Y	MIRROR	COLD	COLD
DATA	1	PARAMETERFILE	COARSE	Y	MIRROR	COLD	COLD
DATA	1	TEMPFILE	COARSE	Y	MIRROR	COLD	COLD
DATA	1	XTRANSPORT	COARSE	Y	MIRROR	COLD	COLD

mktmpl

Purpose

Adds a template to a disk group.

Syntax and Description

```
mktmpl -G diskgroup [--striping {coarse | fine}]
      [--redundancy {high | mirror | unprotected} ] [--primary {hot | cold}]
      [--secondary {hot | cold}] template
```

Table 12-50 Options for the mktmpl Command

Option	Description
-G <i>diskgroup</i>	Name of the disk group.
<i>template</i>	Name of the template to create.
--striping {coarse fine}	Striping specification, either coarse or fine).
--redundancy {high mirror unprotected}	Redundancy specification, either high, mirror, or unprotected.
--primary {hot cold }	Intelligent Data Placement specification for primary extents, either hot or cold region.
--secondary {hot cold }	Intelligent Data Placement specification for secondary extents, either hot or cold region.

For information about adding templates, see ["Adding Templates to a Disk Group"](#) on page 7-17.

Example

The following is an example of the mktmpl command that add the mytemplate template to the data disk group. The new template has the redundancy set to mirror and the striping set to coarse.

Example 12-52 Using mktmpl

```
ASMCMD [+] > mktmpl -G data --redundancy mirror --striping coarse mytemplate
```

rmtmpl

Purpose

Removes a template from a disk group.

Syntax and Description

```
rmtmpl -G diskgroup template
```

Table 12-51 Options for the rmtmpl Command

Option	Description
-G <i>diskgroup</i>	Name of the disk group.
<i>template</i>	Name of the template to delete.

For information about template attributes, see ["Template Attributes"](#) on page 7-15.

Example

The following is an example of the `rmtmpl` command that removes the `mytemplate` template from the `data` disk group.

Example 12–53 Using `rmtmpl`

```
ASMCMD [+] > rmtmpl -G data mytemplate
```

ASMCMD File Access Control Commands

This section describes the ASMCMD Oracle ASM File Access Control commands. For information about Oracle ASM File Access Control, see ["Managing Oracle ASM File Access Control for Disk Groups"](#) on page 4-38.

[Table 12–52](#) provides a summary of ASMCMD Oracle ASM File Access Control commands.

Table 12–52 Summary of ASMCMD File Access Control Commands

Command	Description
<code>chgrp</code>	Changes the user group of a file or list of files.
<code>chmod</code>	Changes permissions of a file or list of files.
<code>chown</code>	Changes the owner of a file or list of files.
<code>groups</code>	List the user groups to which a user belongs.
<code>grpmod</code>	Adds or removes users from an existing user group.
<code>lsgrp</code>	Lists user groups.
<code>lsusr</code>	Lists users in a disk group.
<code>mkgrp</code>	Creates a new user group.
<code>mkusr</code>	Adds a user to disk group.
<code>passwd</code>	Changes the password of a user.
<code>rmgrp</code>	Deletes a user group.
<code>rmusr</code>	Deletes a user from a disk group.

chgrp

Purpose

Changes the user group of a file or list of files.

Syntax and Description

```
chgrp usergroup file [file ...]
```

The syntax options for the `chgrp` command are as follows:

Table 12–53 Options for the `chgrp` Command

Option	Description
<code><i>usergroup</i></code>	Name of the user group.

Table 12–53 (Cont.) Options for the chgrp Command

Option	Description
<i>file</i>	Name of a file.

Only the file owner or the Oracle ASM administrator can use this command. If the user is the file owner, then he must also be either the owner or a member of the group for this command to succeed.

This command accepts a file name or a multiple file names separated by spaces.

Examples

The following are examples of the `chgrp` command that change the Oracle ASM user group of the specified files.

Example 12–54 Using chgrp

```
ASMCMD [+] > chgrp asm_data +data/orcl/controlfile/Current.260.684924747
```

```
ASMCMD [+fra/orcl/archivelog/flashback] > chgrp asm_fra log_7.264.684968167 log_8.265.684972027
```

chmod

Purpose

Changes permissions of a file or list of files.

Syntax and Description

```
chmod mode file [file ...]
```

mode can be one of the following forms:

- { ugo | ug | uo | go | u | g | o | a } { + | - } { r | w | rw }
 - a specifies permissions for all users, u specifies permissions for the owner of the file, g specifies the group permissions, and o specifies permissions for other users.
- { 0 | 4 | 6 } { 0 | 4 | 6 } { 0 | 4 | 6 }
 - The first digit specifies owner permissions, the second digit specifies group permissions, and the third digit specifies other permissions.

The syntax variables for the `chmod` command are described in [Table 12–54](#).

Table 12–54 Options for the chmod Command

Option	Description
6	Read write permissions
4	Read only permissions
0	No permissions
u	Owner permissions, used with r or w
g	Group permissions, used with r or w
o	Other user permissions, used with r or w
a	All user permissions, used with r or w
+	Add a permission, used with r or w

Table 12–54 (Cont.) Options for the chmod Command

Option	Description
-	Removes a permission, used with r or w
r	Read permission
w	Write permission
<i>file</i>	Name of a file

This command accepts a file name or multiple file names separated by spaces. The specified files must be closed.

You can only set file permissions to read-write, read-only, and no permissions. You cannot set file permissions to write-only.

To view the permissions on a file, use the ASMCMD `ls` command with the `--permission` option. See "ls" on page 12-22.

Examples

The following are examples of the `chmod` command that change the permissions of the specified files.

Example 12–55 Using chmod

```
ASMCMD [+fra/orcl/archivelog/flashback] > chmod ug+rw log_7.264.684968167 log_8.265.684972027
```

```
ASMCMD [+fra/orcl/archivelog/flashback] > chmod 640 log_7.264.684968167 log_8.265.684972027
```

```
ASMCMD [+] > ls --permission +fra/orcl/archivelog/flashback
User      Group    Permission  Name
grid      asm_fra  rw-r----- log_7.264.684968167
grid      asm_fra  rw-r----- log_8.265.684972027
```

chown

Purpose

Changes the owner of a file or list of files.

Syntax and Description

```
chown user[:usergroup] file[file...]
```

The syntax options for the `chown` command are as follows:

Table 12–55 Options for the chown Command

Option	Description
<i>user</i>	The name of the user that becomes the new owner.
<i>usergroup</i>	Name of the user group to which the user belongs.
<i>file</i>	Name of a file.

user typically refers to the user that owns the database instance home. Oracle ASM File Access Control uses the operating system (OS) name to identify a database.

This command accepts a file name or multiple file names separated by spaces. The specified files must be closed.

Only the Oracle ASM administrator can use this command.

Examples

The following are examples of the `chown` command that change the owner of the specified files to the `oracle1` operating system user.

Example 12–56 Using `chown`

```
ASMCMD [+fra/orcl/archivelog/flashback] > chown oracle1 log_7.264.684968167 log_8.265.684972027
```

```
ASMCMD [+fra/orcl/archivelog/flashback] > chown oracle1:asm_fra log_9.264.687650269
```

groups

Purpose

Lists all the user groups to which the specified user belongs.

Syntax and Description

```
groups diskgroup user
```

The syntax variables for the `groups` command are as follows:

Table 12–56 Options for the `groups` Command

Option	Description
<i>diskgroup</i>	Name of the disk group to which the user belongs.
<i>user</i>	Name of the user.

Examples

The following is an example of the `groups` command that displays the user groups of the `data` disk group to which the `oracle1` user belongs.

Example 12–57 Using `groups`

```
ASMCMD [+] > groups data oracle1
asm_data
```

grpmod

Purpose

Adds or removes operating system (OS) users to and from an existing Oracle ASM user group.

Syntax and Description

```
grpmod { --add | --delete } diskgroup usergroup user [user...]
```

Table 12–57 Options for the *grpmod* Command

Option	Description
<code>--add</code>	Specifies to add users to the user group.
<code>--delete</code>	Specifies to delete users from the user group.
<i>diskgroup</i>	Name of the disk group to which the user group belongs.
<i>usergroup</i>	Name of the user group.
<i>user</i>	Name of the user to add or remove from the user group.

Only the owner of the user group can use this command. The command requires the SYSASM privilege to run.

This command accepts an operating system user name or multiple user names separated by spaces. The operating system users are typically owners of a database instance home.

Examples

The following are examples of the *grpmod* command. The first example adds the *oracle1* and *oracle2* users to the *asm_fra* user group of the *fra* disk group. The second example removes the *oracle2* user from the *asm_data* user group of the *data* disk group.

Example 12–58 Using *grpmod*

```
ASMCMD [+] > grpmod --add fra asm_fra oracle1 oracle2
```

```
ASMCMD [+] > grpmod --delete data asm_data oracle2
```

lsgrp

Purpose

Lists all Oracle ASM user groups or only groups that match a specified pattern.

Syntax and Description

```
lsgrp [-Ha] [-G diskgroup] [pattern]
```

Table 12–58 Options for the *lsgrp* command

Option	Description
<code>-H</code>	Suppresses column headings.
<code>-a</code>	Lists all columns.
<code>-G <i>diskgroup</i></code>	Limits the results to the specified disk group name.
<i>pattern</i>	Displays the user groups that match the pattern expression.

Examples

The following are examples of the *lsgrp* command. The first example displays a subset of information about the user groups whose name matches the *asm%* pattern. The second example displays all information about all the user groups.

Example 12–59 Using lsgrp

```

ASMCMD [+] > lsgrp asm%
DG_Name Grp_Name Owner
FRA     asm_fra  grid
DATA    asm_data grid

ASMCMD [+] > lsgrp -a
DG_Name Grp_Name Owner Members
FRA     asm_fra  grid  oracle1
DATA    asm_data grid  oracle1 oracle2
    
```

lsusr

Purpose

Lists Oracle ASM users in a disk group.

Syntax and Description

```
lsusr [-Ha] [ -G diskgroup] [pattern]
```

Table 12–59 Options for the lsusr command

Option	Description
-a	List all users and the disk groups to which the users belongs.
-H	Suppresses column headings.
-G <i>diskgroup</i>	Limits the results to the specified disk group name.
<i>pattern</i>	Displays the users that match the pattern expression.

Examples

The following is an example of the `lsusr` command. The example lists users in the data disk group and also shows the operating system user Id assigned to the user.

Example 12–60 Using lsusr

```

ASMCMD [+] > lsusr -G data
User_Num OS_ID OS_Name
3        1001  grid
1        1021  oracle1
2        1022  oracle2
    
```

mkgrp

Purpose

Creates a new Oracle ASM user group.

Syntax and Description

```
mkgrp diskgroup usergroup [user] [user...]
```

Table 12–60 Options for the *mkgrp* Command

Option	Description
<i>diskgroup</i>	Name of the disk group to which the user group is added.
<i>usergroup</i>	Name of the user group to add. 30 is the maximum number of characters.
<i>user</i>	Name of the database user to add to the user group.

You can optionally specify a list of users to be included as members of the new user group.

Examples

The following is an example of the *mkgrp* command. This example creates the *asm_data* user group in the *data* disk group and adds the *oracle1* and *oracle2* users to the user group.

Example 12–61 Using *mkgrp*

```
ASMCMD [+] > mkgrp data asm_data oracle1 oracle2
```

mkusr

Purpose

Adds an operating system (OS) user to a disk group.

Syntax and Description

```
mkusr diskgroup user
```

The syntax variables for the *mkusr* command are as follows:

Table 12–61 Options for the *mkusr* command

Option	Description
<i>diskgroup</i>	Specifies the name of the disk group to which the user is to be added.
<i>user</i>	Name of the user that you want to add.

The user to be added must be a valid operating system user. Only a user authenticated as *SYSASM* can execute this command.

Examples

The following are examples of the *mkusr* command. The first example adds the *oracle1* user to the *data* disk group. The second example adds the *oracle2* user to the *fra* disk group.

Example 12–62 Using *mkusr*

```
ASMCMD [+] > mkusr data oracle1
```

```
ASMCMD [+] > mkusr fra oracle2
```


passwd

Purpose

Changes the password of a user.

Syntax and Description

```
passwd user
```

The syntax variables for the `passwd` command are as follows:

Table 12–62 Options for the `passwd` command

Option	Description
<i>user</i>	Name of the user.

An error is raised if the user does not exist in the Oracle ASM password file. The user is first prompted for the current password, then the new password. The command requires the SYSASM privilege to run.

Examples

The following is an example of the `passwd` command that changes the password of the `oracle2` user.

Example 12–63 Using `passwd`

```
ASMCMD [+] > passwd oracle2
Enter old password (optional):
Enter new password: *****
```

rmgrp

Purpose

Removes a user group from a disk group.

Syntax and Description

```
rmgrp diskgroup usergroup
```

Table 12–63 Options for the `rmgrp` Command

Option	Description
<i>diskgroup</i>	Name of the disk group to which the user group belongs.
<i>usergroup</i>	Name of the user group to delete.

Removing a group might leave some files without a valid group. To ensure that those files have a valid group, explicitly update those files to a valid group. See "[chgrp](#)" on page 12-50.

The command must be run by the owner of the group and also requires the SYSASM privilege to run.

Examples

The following is an example of the `rmgrp` command that removes the `asm_data` user group from the `data` disk group.

Example 12–64 Using `rmgrp`

```
ASMCMD [+] > rmgrp data asm_data
```

rmusr

Purpose

Deletes an operating system (OS) user from a disk group.

Syntax and Description

```
rmusr [-r] diskgroup user
```

Table 12–64 Options for the `rmusr` command

Option	Description
<code>-r</code>	Removes all files in the disk group that the user owns at the same time that the user is removed.
<i>diskgroup</i>	Specifies the name of the disk group from which the user is to be deleted.
<i>user</i>	Name of the user that you want to delete.

Only a user authenticated as `SYSASM` can execute this command.

Examples

The following is an example of the `rmusr` command that removes the `oracle2` user from the `data` disk group.

Example 12–65 Using `rmusr`

```
ASMCMD [+] > rmusr data oracle2
```

ASMCMD Volume Management Commands

This section describes the ASMCMD volume management commands.

For information about Oracle ASM Dynamic Volume Manager (Oracle ADVM), see ["Overview of Oracle ASM Dynamic Volume Manager"](#) on page 5-10.

[Table 12–65](#) provides a summary of the Oracle ADVM volume management commands. To successfully execute these commands, the local Oracle ASM instance must be running and the disk group required by this command must have been created and mounted in the Oracle ASM instance.

Table 12–65 Summary of ASMCMD Volume Management Commands

Command	Description
<code>volcreate</code>	Creates an Oracle ADVM volume in the specified disk group.
<code>voldelete</code>	Deletes an Oracle ADVM volume.
<code>voldisable</code>	Disables Oracle ADVM volumes in mounted disk groups.
<code>volenable</code>	Enables Oracle ADVM volumes in mounted disk groups.
<code>volinfo</code>	Displays information about Oracle ADVM volumes.
<code>volresize</code>	Resizes an Oracle ADVM volume.
<code>volset</code>	Sets attributes of an Oracle ADVM volume in mounted disk groups.
<code>volstat</code>	Reports volume I/O statistics.

volcreate

Purpose

Creates an Oracle ADVM volume in the specified disk group.

Syntax and Description

```
volcreate -G diskgroup -s size
      [ --column number ] [ --width stripe_width ]
      [ --redundancy {high|mirror|unprotected} ]
      [ --primary {hot|cold} ] [ --secondary {hot|cold} ] volume
```

Table 12–66 describes the options for the `volcreate` command.

Table 12–66 Options for the volcreate command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the disk group containing the volume.
<code>-s <i>size</i></code>	Size of the volume to be created in units of K, M, G, or T. The unit designation must be appended to the number specified. No space is allowed. For example: 20G
<code>--column <i>number</i></code>	Number of columns in a stripe set. Values range from 1 to 8. The default value is 4.
<code>--width <i>stripe</i></code>	Stripe width of a volume. The value can range from 4 KB to 1 MB, at power-of-two intervals, with a default of 128 KB.
<code>--redundancy {high mirror unprotected}</code>	Redundancy of the Oracle ADVM volume which can be specified for normal redundancy disk groups. The range of values are as follows: <code>unprotected</code> for non-mirrored redundancy, <code>mirror</code> for double-mirrored redundancy, or <code>high</code> for triple-mirrored redundancy. If redundancy is not specified, the setting defaults to the redundancy level of the disk group.
<code>--primary {hot cold}</code>	Intelligent Data Placement specification for primary extents, either <code>hot</code> or <code>cold</code> region.
<code>--secondary {hot cold}</code>	Intelligent Data Placement specification for secondary extents, either <code>hot</code> or <code>cold</code> region.

Table 12–66 (Cont.) Options for the volcreate command

Option	Description
<i>volume</i>	Name of the volume to be created. Can be a maximum of 11 alphanumeric characters; dashes are not allowed. The first character must be alphabetic.

When creating an Oracle ADVM volume, a volume device name is created with a unique Oracle ADVM persistent disk group number that is concatenated to the end of the volume name. The unique number can be one to three digits.

On Linux, the volume device name is in the format *volume_name-nnn*. On Windows the volume device name is in the format *asm-volume_name-nnn*. The volume device file can be used as any other disk or logical volume to mount file systems or for applications to use directly. For information on mounting the volume device file, see ["Creating an Oracle ACFS File System"](#) on page 13-1.

You can determine the volume device name with the `volinfo` command, described in ["volinfo"](#) on page 12-62.

A successful volume creation automatically enables the volume device.

The volume device file can be used as any other disk or logical volume to mount file systems or for applications to use directly.

Examples

The following is an example of the `volcreate` command that creates `volume1` in the data disk group with the size set to 10 gigabytes.

Example 12–66 Using volcreate

```
ASMCMD [+] > volcreate -G data -s 10G --width 64K --column 8 volume1
```

```
ASMCMD [+] > volinfo -G data volume1
Diskgroup Name: DATA
```

```
Volume Name: VOLUME1
Volume Device: /dev/asm/volume1-123
State: ENABLED
Size (MB): 10240
Resize Unit (MB): 512
Redundancy: MIRROR
Stripe Columns: 8
Stripe Width (K): 64
Usage:
Mountpath:
```

voldelete

Purpose

Deletes an Oracle ADVM volume.

Syntax and Description

```
voldelete -G diskgroup volume
```

[Table 12–67](#) describes the options for the `voldelete` command.

Table 12–67 Options for the voldelete command

Option	Description
<code>-G diskgroup</code>	Name of the disk group containing the volume.
<code>volume</code>	Name of the volume to be operated on.

To successfully execute this command, the local Oracle ASM instance must be running and the disk group required by this command must be mounted in the Oracle ASM instance. Before deleting a volume, you must ensure that there are no active file systems associated with the volume.

Examples

The following is an example of the `voldelete` command that deletes `volume1` from the `data` disk group.

Example 12–67 Using voldelete

```
ASMCMD [+] > voldelete -G data volume1
```

voldisable

Purpose

Disables Oracle ADVM volumes in mounted disk groups and removes the volume device on the local node.

Syntax and Description

```
voldisable { -a | -G diskgroup -a | -G diskgroup volume }
```

[Table 12–68](#) describes the options for the `voldisable` command.

Table 12–68 Options for the voldisable command

Option	Description
<code>-a</code>	When used without a disk group name, specifies all volumes within all disk groups. When used with a disk group name (<code>-G diskgroup -a</code>), specifies all volumes within that disk group.
<code>-G diskgroup</code>	Name of the disk group containing the volume.
<code>volume</code>	Name of the volume to be operated on. Can be maximum of 30 alphanumeric characters. The first character must be alphabetic.

You can disable volumes before shutting down an Oracle ASM instance or dismounting a disk group to verify that the operations can be accomplished normally without including a force option due to open volume files. Disabling a volume also prevents any subsequent opens on the volume or device file as it no longer exists.

Before disabling a volume, you must ensure that there are no active file systems associated with the volume. You must first dismount the Oracle ACFS file system before disabling the volume. See ["Dismounting an Oracle ACFS File System"](#) on page 13-5.

You can delete a volume without first disabling the volume.

Examples

The following is an example of the `voldisable` command that disables `volume1` in the `data` disk group.

Example 12–68 Using voldisable

```
ASMCMD [+] > voldisable -G data volume1
```

volenable

Purpose

Enables Oracle ADVM volumes in mounted disk groups.

Syntax and Description

```
volenable { -a | -G diskgroup -a | -G diskgroup volume }
```

Table 12–69 describes the options for the `volenable` command.

Table 12–69 Options for the volenable command

Option	Description
-a	When used without a disk group name, specifies all volumes within all disk groups. When used with a disk group name (<code>-G <i>diskgroup</i> -a</code>), specifies all volumes within that disk group.
-G <i>diskgroup</i>	Name of the disk group containing the volume.
<i>volume</i>	Name of the volume to be operated on.

A volume is enabled when it is created.

Examples

The following is an example of the `volenable` command that enables `volume1` in the `data` disk group.

Example 12–69 Using volenable

```
ASMCMD [+] > volenable -G data volume1
```

volinfo

Purpose

Displays information about Oracle ADVM volumes.

Syntax and Description

```
volinfo { -a | -G diskgroup -a | -G diskgroup volume }  
volinfo [--show_diskgroup | --show_volume] volumedev }
```

Table 12–70 describes the options for the `volinfo` command.

Table 12–70 Options for the volinfo command

Option	Description
-a	When used without a disk group name, specifies all volumes within all disk groups. When used with a disk group name (<i>-G diskgroup -a</i>), specifies all volumes within that disk group.
<i>-G diskgroup</i>	Name of the disk group containing the volume.
<i>volume</i>	Name of the volume.
--show_diskgroup	Returns only the disk group name. A volume device name is required.
--show_volume	Returns only the volume name. A volume device name is required.
<i>volumedevic</i>	Name of the volume device.

Examples

The following example displays information about the `volume1` volume in the `data` disk group.

Example 12–70 Using volinfo

```
ASMCMD [+] > volinfo -G data volume1
Diskgroup Name: DATA
  Volume Name: VOLUME1
  Volume Device: /dev/asm/volume1-123
  State: ENABLED
  Size (MB): 10240
  Resize Unit (MB): 512
  Redundancy: MIRROR
  Stripe Columns: 8
  Stripe Width (K): 64
  Usage: ACFS
  Mountpath: /u01/app/acfsmounts/acfs1
```

volresize

Purpose

Resizes an Oracle ADVM volume.

Syntax and Description

```
volresize -G diskgroup -s size [ -f ] volume
```

Table 12–71 describes the options for the `volresize` command.

Table 12–71 Options for the volresize command

Option	Description
<i>-G diskgroup</i>	Name of the disk group containing the volume.
-f	Force the shrinking of a volume that is not an Oracle ACFS volume to suppress the warning message.
<i>volume</i>	Name of the volume to be operated on.
<i>-s size</i>	New size of the volume in units of M, G, or T.

If the volume is not an Oracle ACFS volume and the new size is smaller than current, you are warned of possible data corruption. Unless the `-f` (force) option is specified, you are prompted whether to continue with the operation.

If there is an Oracle ACFS file system on the volume you cannot resize the volume with the `volresize` command. You must use the `acfsutil size` command, which also resizes the volume and file system. For information, see "[acfsutil size](#)" on page 13-29.

Examples

The following is an example of the `volresize` command that resizes `volume1` in the data disk group to 20 gigabytes.

Example 12-71 Using volresize

```
ASMCMD [+] > volresize -G data -s 20G volume1
```

volset

Purpose

Sets attributes of an Oracle ADVM volume in mounted disk groups.

Syntax and Description

```
volset -G diskgroup [ --usagestring string ]
      [ --mountpath mount_path ]
      [ --primary {hot|cold} ] [ --secondary {hot|cold} ] volume
```

[Table 12-72](#) describes the options for the `volset` command.

Table 12-72 Options for the volset command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the disk group containing the volume.
<code>--usagestring <i>string</i></code>	Optional usage string to tag a volume which can be up to 30 characters. This string is set to ACFS when the volume is attached to a file system and should not be changed.
<code>--mountpath <i>mount_path</i></code>	Optional string to tag a volume with its mount path string which can be up to 1024 characters. This string is for user information only.
<code>--primary {hot cold}</code>	Intelligent Data Placement specification for primary extents, either <code>hot</code> or <code>cold</code> region.
<code>--secondary {hot cold}</code>	Intelligent Data Placement specification for secondary extents, either <code>hot</code> or <code>cold</code> region.
<code><i>volume</i></code>	Name of the volume to be operated on.

Examples

The following is an example of the `volset` command that sets the usage string for a volume that is not attached to a file system.

Example 12-72 Using volset

```
ASMCMD [+] > volset -G data --usagestring 'no file system attached' volume1
```


volstat

Purpose

Reports I/O statistics for Oracle ADVM volumes.

Syntax and Description

```
volstat [-G diskgroup] [volume]
```

[Table 12-73](#) describes the options for the `volstat` command.

Table 12-73 Options for the `volstat` command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the mounted disk group containing the volume.
<code><i>volume</i></code>	Name of the volume to be operated on.

The following apply when using the `volstat` command.

- If the disk group is not specified and the volume name is specified, all mounted disk groups are searched for the specified volume name.
- If the disk group name is specified and the volume name is omitted, all volumes are displayed for the named disk group.
- If both the disk group name and the volume name are omitted, all volumes on all disk groups are displayed.

Examples

The following is an example of the `volstat` command that displays information about volumes in the data disk group.

Example 12-73 Using `volstat`

```
ASMCMD [+] > volstat -G data
DISKGROUP NUMBER / NAME: 1 / DATA
-----
VOLUME_NAME
  READS          BYTES_READ    READ_TIME     READ_ERRS
  WRITES         BYTES_WRITTEN WRITE_TIME     WRITE_ERRS
-----
VOLUME1
  10085          2290573312   22923         0
  1382           5309440     1482          0
```

Oracle ACFS Command-Line Tools

Command-line tools for managing and implementing Oracle Automatic Storage Management Cluster File System (Oracle ACFS), include the following:

- [Basic Steps to Manage an Oracle ACFS](#)
- [Oracle ACFS Command-line Tools for Linux and UNIX Environments](#)
- [Oracle ACFS Command-line Tools for Windows Environments](#)
- [Oracle ACFS Command-Line Tools for All Environments](#)

Basic Steps to Manage an Oracle ACFS

This section provides an outline of the basic steps when managing an Oracle ACFS using command-line utilities.

The examples in this section show operating system commands that are run in a Linux environment system. ASMCMD commands are used to manage the volumes, but SQL*Plus, Oracle ASM Configuration Assistant (ASMCA), and Oracle Enterprise Manager can also be used to manage volumes.

This section contains these topics:

- [Creating an Oracle ACFS File System](#)
- [Accessing an Oracle ACFS File System on a Different Node in the Cluster](#)
- [Managing Oracle ACFS Snapshots](#)
- [Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#)
- [Removing an Oracle ACFS File System and a Volume](#)

Creating an Oracle ACFS File System

To create and verify a file system, perform the following steps:

1. Create an Oracle ASM volume in a mounted disk group with the ASMCMD `volcreate` command.

The compatibility parameters `COMPATIBLE.ASM` and `COMPATIBLE.ADVM` must be set to 11.2 or higher for the disk group. See "[Disk Group Compatibility Attributes](#)" on page 4-33.

Start ASMCMD connected to the Oracle ASM instance. You must be a user in the OSASM operating system group. See "[About Privileges for Oracle ASM](#)" on page 3-23.

When configuring Oracle ADVM volume devices within a disk group, Oracle recommends assigning the Oracle grid infrastructure user and Oracle ASM administrator roles to users who have root privileges.

To create a volume:

```
ASMCMD [+] > volcreate -G data -s 10G volume1
```

The volume name must be less than or equal to eleven alphanumeric characters, starting with an alphabetic character.

When creating an Oracle ASM volume, a volume device name is created that includes a unique Oracle ADVM persistent disk group number. The volume device file can be used in the same manner as any other disk or logical volume to mount file systems or for applications to use directly.

For information about the `volcreate` command, see "[volcreate](#)" on page 12-59.

2. Determine the device name of the volume that was created.

You can determine the volume device name with the `ASMCMD volinfo` command or from the `VOLUME_DEVICE` column in the `V$ASM_VOLUME` view.

For example:

```
ASMCMD [+] > volinfo -G data volume1
Diskgroup Name: DATA
```

```
Volume Name: VOLUME1
Volume Device: /dev/asm/volume1-123
State: ENABLED
```

...

```
SQL> SELECT volume_name, volume_device FROM V$ASM_VOLUME
WHERE volume_name = 'VOLUME1';
```

VOLUME_NAME	VOLUME_DEVICE
VOLUME1	/dev/asm/volume1-123

For information about the `volinfo` command, see "[volinfo](#)" on page 12-62.

See Also: *Oracle Database Reference* for information about the `V$ASM_VOLUME` view

3. Create a file system with the Oracle ACFS `mkfs` command.

Create a file system using an existing volume device.

For example:

```
$ /sbin/mkfs -t acfs /dev/asm/volume1-123
```

```
mkfs.acfs: version = 11.2.0.1.0.0
mkfs.acfs: on-disk version = 39.0
mkfs.acfs: volume = /dev/asm/volume1-123
mkfs.acfs: volume size = 10737418240
mkfs.acfs: Format complete.
```

See "[mkfs](#)" on page 13-9 (Linux or UNIX) or "[acfsformat](#)" on page 13-16 (Windows). The `root` privilege is not required. The ownership of the volume device file dictates who can run this command.

4. Register the file system with the `acfsutil registry` command.

For example:

```
$ /sbin/acfsutil registry -a /dev/asm/volume1-123 /u01/app/acfsmounts/myacfs
acfsutil registry: mount point /u01/app/acfsmounts/myacfs successfully added
to Oracle Registry
```

See "[acfsutil registry](#)" on page 13-27. The `root` or `asmadmin` privileges are required to modify the registry. The Windows Administrator privilege is equivalent to the `root` privilege on Linux.

Registering a file system is optional. After registering an Oracle ACFS file system in the cluster mount registry, the file system is mounted automatically on each cluster member listed in the registry entry during the next registry check action. This automatic process runs every 30 seconds and eliminates the need to manually mount the file system on each member of the cluster.

Registering an Oracle ACFS file system also causes the file system to be mounted automatically whenever Oracle Clusterware or the system is restarted.

Note: In an Oracle grid infrastructure Clusterware configuration, you can run `srvctl add filesystem` to automount a file system; this method is required when an Oracle Database home is installed on an Oracle ACFS file system. However, that file system should not be added to the registry. For information about Server Control Utility (SRVCTL), see *Oracle Real Application Clusters Administration and Deployment Guide*.

For more information, see "[About the Oracle ACFS Mount Registry](#)" on page 5-6.

Note: A file system is not automatically mounted for an Oracle Restart configuration, which is a single-instance (non-clustered) environment.

5. Mount the file system with the Oracle ACFS `mount` command. You can mount a file system before or after registering the file system. If the file system has been registered, you can wait for the file system to be mounted automatically.

For example:

```
# /bin/mount -t acfs /dev/asm/volume1-123 /u01/app/acfsmounts/myacfs
```

See "[mount](#)" on page 13-10 (Linux or UNIX) or "[acfsmountvol](#)" on page 13-17 (Windows). The `root` privilege is required run the `mount` command and the Windows Administrator privilege is required to run the `acfsmountvol` command.

After the file system has been mounted, you need to ensure that the permissions are set to allow access to the file system for the appropriate users. For example:

```
# chown -R oracle:dba /u01/app/acfsmounts/myacfs
```

6. Create a test file in the file system.

The user that creates the test file should be a user that is intended to access the file system. This test ensures that the appropriate user can write to the file system.

For example:

```
$ echo "Oracle ACFS File System" > /u01/app/acfsmounts/myacfs/myfile
```

7. List the contents of the test file that was created in the file system.

For example:

```
$ cat /u01/app/acfsmounts/myacfs/myfile
Oracle ACFS File System
```

Accessing an Oracle ACFS File System on a Different Node in the Cluster

If the node is part of a cluster system, perform the following steps on node 2 to view the test file you created on node 1.

Note: If the file system has been registered with the Oracle ACFS mount registry, you can skip steps 1 to 3.

1. Enable the volume that was previously created and enabled on node 1.

Start ASMCMD connected to the Oracle ASM instance. You must be a user in the OSASM operating system group. See "[About Privileges for Oracle ASM](#)" on page 3-23.

For example:

```
ASMCMD [+] > volenable -G data volume1
```

See "[volenable](#)" on page 12-62.

2. View information about the volume that you created on node 1.

For example:

```
ASMCMD [+] > volinfo -G data volume1
```

See "[volinfo](#)" on page 12-62.

3. Mount the file system using the Oracle ACFS mount command.

For example:

```
# /bin/mount -t acfs /dev/asm/volume1-123 /u01/app/acfsmounts/myacfs
```

See "[mount](#)" on page 13-10 (Linux or UNIX) or "[acfsmountvol](#)" on page 13-17 (Windows). The `root` privilege is required run the `mount` command and the Windows Administrator privilege is required to run the `acfsmountvol` command.

After the file system has been mounted, you need to ensure that the permissions are set to allow access for the appropriate users.

4. List the contents of the test file you previously created on the file system.

For example:

```
$ cat /u01/app/acfsmounts/myacfs/myfile
Oracle ACFS File System
```

The contents should match the file created previously on node 1.

Managing Oracle ACFS Snapshots

To create and verify a snapshot on node 1:

1. Create snapshot of the new file system created on node 1.

For example:

```
$ /sbin/acfsutil snap create mysnapshot_20090725 /u01/app/acfsmounts/myacfs
```

See ["acfsutil snap create"](#) on page 13-30.

2. Update the test file in the file system so that it is different than the snapshot.

For example:

```
$ echo "Modifying a file in Oracle ACFS File System" >
/u01/app/acfsmounts/myacfs/myfile
```

3. List the contents of the test file and the snapshot view of the test file.

For example:

```
$ cat /u01/app/acfsmounts/myacfs/myfile

$ cat /u01/app/acfsmounts/myacfs/.ACFS/snaps/mysnapshot_20090725/myfile
```

The contents of the test file and snapshot should be different. If node 1 is in a cluster, then you can perform the same list operation on node 2.

To manage snapshots with Oracle Enterprise Manager, see ["Managing Oracle ACFS Snapshots with Oracle Enterprise Manager"](#) on page 10-5.

Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems

This sections discusses the operations to deregister or dismount a file system and disable a volume. This section contains these topics:

- [Deregistering an Oracle ACFS File System](#)
- [Dismounting an Oracle ACFS File System](#)
- [Disabling a Volume](#)

Deregistering an Oracle ACFS File System

You can deregister an Oracle ACFS file system if you do not want the file system to be automatically mounted.

For example:

```
$ /sbin/acfsutil registry -d /u01/app/acfsmounts/myacfs
```

If you deregister a file system, then you must explicitly mount the file system after Oracle Clusterware or the system is restarted.

For more information about the registry, see ["About the Oracle ACFS Mount Registry"](#) on page 5-6. For information about `acfsutil registry`, see ["acfsutil registry"](#) on page 13-27.

Dismounting an Oracle ACFS File System

You can dismount a file system without deregistering the file system or disabling the volume on which the file system is mounted.

For example, you can dismount a file system and run `fsck` to check the file system.

```
# /bin/umount /u01/app/acfsmounts/myacfs
# /sbin/fsck -a -v -y -t acfs /dev/asm/volume1-123
```

After you dismount a file system, you must explicitly mount the file system.

Use `umount` on Linux systems or `acfsdismount` on Windows systems. For information about the commands to dismount a file system, see "[umount](#)" on page 13-12 or "[acfsdismount](#)" on page 13-15.

Use `fsck` on Linux systems or `acfschkdsk` on Windows systems to check a file system. For information about the commands to check a file system, see "[fsck](#)" on page 13-7 or "[acfschkdsk](#)" on page 13-13.

Disabling a Volume

To disable a volume, you must first dismount the file system on which the volume is mounted.

For example:

```
# /bin/umount /u01/app/acfsmounts/myacfs
```

After a file system is dismounted, you can disable the volume and remove the volume device file.

For example:

```
ASMCMD> voldisable -G data volume1
```

Dismounting the file system and disabling a volume does not destroy data in the file system. You can enable the volume and mount the file system to access the existing data. For information about `voldisable` and `volenable`, see "[voldisable](#)" on page 12-61 and "[volenable](#)" on page 12-62.

Removing an Oracle ACFS File System and a Volume

To permanently remove a volume and Oracle ACFS file system, perform the following steps. These steps destroy the data in the file system.

1. Deregister the file system with `acfsutil registry -d`.

For example:

```
$ /sbin/acfsutil registry -d /oracle/acfsmounts/acfs1
acfsutil registry: successfully removed ACFS mount point
/oracle/acfsmounts/acfs1 from Oracle Registry
```

For information about running `acfsutil registry`, see "[acfsutil registry](#)" on page 13-27.

2. Dismount the file system.

For example:

```
# /bin/umount /oracle/acfsmounts/acfs1
```

You must dismount the file system on all nodes of a cluster.

Use `umount` on Linux systems or `acfsdismount` on Windows systems. For information about running `umount` or `acfsdismount`, see "[umount](#)" on page 13-12 or "[acfsdismount](#)" on page 13-15.

3. Remove the file system with `acfsutil rmfs`.

If you were not planning to remove the volume in a later step, this step is necessary to remove the file system. Otherwise, the file system is removed when the volume is deleted.

For example:

```
$ /sbin/acfsutil rmfs /dev/asm/volume1-123
```

For information about running `acfsutil rmfs`, see "[acfsutil rmfs](#)" on page 13-28.

4. Optionally you can disable the volume with the ASMCMD `voldisable` command.

For example:

```
ASMCMD> voldisable -G data volume1
```

For information about running `voldisable`, see "[voldisable](#)" on page 12-61.

5. Delete the volume with the ASMCMD `voldelete` command.

For example:

```
ASMCMD> voldelete -G data volume1
```

For information about running `voldelete`, see "[voldelete](#)" on page 12-60.

Oracle ACFS Command-line Tools for Linux and UNIX Environments

[Table 13–1](#) contains a summary of the Oracle ACFS commands for Linux and UNIX.

Table 13–1 Summary of ACFS Commands for Linux and UNIX

Command	Description
fsck	Checks and repairs an Oracle ACFS file system.
mkfs	Creates an Oracle ACFS file system.
mount	Mounts an Oracle ACFS file system.
umount	Dismounts an Oracle ACFS file system.

The commands in [Table 13–1](#) have been extended with additional options to support Oracle ACFS. All other Linux and UNIX file system commands operate without change for Oracle ACFS.

For example, Oracle ACFS adds a set of Oracle ACFS-specific mount options to those provided with the base operating system platform. You should review both the mount options for the Linux and UNIX platforms in addition to the Oracle ACFS-specific options for the complete set of file system mount options.

File systems on Oracle ADVM volumes that are not Oracle ACFS file systems, such as `ext3`, are managed with the same Linux commands that are listed in [Table 13–1](#) using the file-specific options for the type of file system. You can refer to the `man` pages for options available for the Linux commands in [Table 13–1](#).

fsck

Purpose

Checks and repairs an Oracle ACFS file system.

Syntax and Description

```
fscck [-h]
fscck [-a|-f] [-v] -t acfs [-n|-y] volume_device
```

Table 13–2 contains the options available with the `fscck` command.

Table 13–2 Options for the `fscck` command

Option	Description
-a	Specifies to automatically fix the file system.
-f	Forces the file system into mountable state without completing a file system check or fix.
-v	Specifies verbose mode. The progress is displayed as the operation occurs.
-h	Displays the usage help text and exits.
-n	Answers no to any prompts.
-y	Answers yes to any prompts.
<i>volume_device</i>	Specifies the primary Oracle ADVM volume device.

`fscck` checks and repairs an existing Oracle ACFS. This command can only be run on a dismounted file system. `root` privileges are required to run `fscck`. The Oracle ACFS driver must be loaded for `fscck` to work.

By default, `fscck` only checks for and reports any errors. The `-a` flag must be specified to instruct `fscck` to fix errors in the file system.

In a few cases, `fscck` prompts for questions before proceeding to check a file system. These cases include:

- If `fscck` detects that another `fscck` is in progress on the file system
- If `fscck` detects that the Oracle ACFS driver is not loaded
- If the file system does not appear to be Oracle ACFS

In checking mode, `fscck` also prompts if there are transaction logs that have not been processed completely due to an unclean shutdown. To run in a non-interactive mode, include either the `-y` or `-n` options to answer yes or no to any questions.

`fscck` creates working files before it checks a file system. These working files are created in `/usr/tmp` if space is available. `/tmp` is used if `/usr/tmp` does not exist. If insufficient space is available in the `tmp` directory, `fscck` attempts to write to the current working directory. The files that `fscck` creates are roughly the size of the file system being checked divided by 32K. At most two such files are allocated. For example, a 2 GB file system being checked causes `fscck` to generate one or two 64K working files in the `/usr/tmp` directory. These files are deleted after `fscck` has finished.

In the event that `fscck` finds a file or directory in the file system for which it cannot determine its name or intended location (possibly due to a corruption in its parent directory), it places this object in the `/lost+found` directory when `fscck` is run in fix mode. For security reasons only the `root` user on Linux is able to read files in `/lost+found`. If the administrator can later determine the original name and location of the file based on its contents, the file can be moved or copied into its intended location.

The file names in the `/lost+found` directory are in the following formats:

```
parent.id.file.id.time-in-sec-since-1970
parent.id.dir.id.time-in-sec-since-1970
```

The *id* fields are the internal Oracle ACFS numeric identifiers for each file and directory in the file system.

You can use `acfsutil info id id mount_point` to attempt to determine the directory associated with `parent.id`. This directory is assumed to be where the deleted object originated. For information about `acfsutil info`, see "acfsutil info" on page 13-25.

If the parent directory is not known, the parent *id* field is set to UNKNOWN.

Note: It is not possible to see the contents of the `/lost+found` directory from a snapshot.

Examples

The following example shows how to check and repair an Oracle ACFS file system.

Example 13-1 Using fsck

```
# /sbin/fsck -a -v -y -t acfs /dev/asm/volume1-123
```

mkfs

Purpose

Creates an Oracle ACFS file system.

Syntax and Description

```
mkfs [-h]
mkfs [-v] [-f] -t acfs [-b blocksize] [-n name] volume_device [blocks]
```

Table 13-3 contains the options available with the `mkfs` command.

Table 13-3 Options for the `mkfs` command

Option	Description
-t <i>acfs</i>	Specifies the type of file system on Linux. <i>acfs</i> designates the Oracle ACFS type.
-v	Specifies verbose mode. The progress is displayed as the operation occurs.
-n <i>name</i>	Specifies the name for the file system. A name can be a maximum of 64 characters. <code>acfsutil info fs</code> returns the name if one was specified.
-f	Specifies the force option. This action creates the file system even if there is an existing Oracle ACFS on the volume device, although only if the file system is dismounted. This option overwrites structures on the original file system and should be used with caution.
-h	Displays the usage help text and exits.
-b <i>blocksize</i>	The default block size is 4K and this is the only size supported in 11g Release 2 (11.2).

Table 13–3 (Cont.) Options for the mkfs command

Option	Description
<i>blocks</i>	Specifies the number of blocks that the file system should consume on the named device. The quantity specified can be in units of K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes). If a unit is not specified, the default is bytes. If number of blocks specified is not a multiple of the block size, than the value is rounded up to the closest multiple. If this option is not specified, the entire device is consumed.
<i>volume_device</i>	Specifies an Oracle ADVM device file that is to be formatted.

`mkfs` is used to create the on-disk structure needed for Oracle ACFS file system to be mounted. The `mkfs` command is the traditional UNIX command used to build a file system. After `mkfs` executes successfully, the `USAGE` column in the `V$ASM_VOLUME` view displays `ACFS`. `root` privilege is not required. The ownership of the volume device file dictates who can run this command. The minimum file system size is 200 MB. The Oracle ACFS driver must be loaded for `mkfs` to work.

Examples

Before creating an Oracle ACFS file system, first determine which Oracle ADVM volume devices are available. You can use the `ASMCMD volinfo` command to display information about the volumes and volume devices.

```
ASMCMD [+] > volinfo -a
...
      Volume Name: VOLUME1
      Volume Device: /dev/asm/volume1-123
      State: ENABLED
...
```

See "[volinfo](#)" on page 12-62.

Next create an Oracle ACFS file system on the volume device file.

Example 13–2 Using mkfs

```
$ /sbin/mkfs -t acfs /dev/asm/volume1-123
```

mount

Purpose

Mounts an Oracle ACFS file system.

Syntax and Description

```
mount [-h]
mount [-v] -t acfs [-o options] volume_device dir
mount
```

[Table 13–4](#) contains the options available with the `mount` command.

Table 13–4 Options for the mount command

Option	Description
<code>-h</code>	Displays the usage help text and exits.

Table 13–4 (Cont.) Options for the mount command

Option	Description
<code>-t acfs</code>	Specifies the type of file system on Linux. <code>acfs</code> designates the Oracle ACFS type.
<code>-v</code>	Specifies verbose mode. The progress is displayed as the operation occurs.
<code>-o</code>	Options are specified with the <code>-o</code> flag followed by a comma separated string of options. The following options are available: <ul style="list-style-type: none"> ▪ <code>all</code> Reads the Oracle ACFS mount registry created with the <code>acfsutil registry</code> command and mounts the file systems in it. A <code>mount -t acfs -o all</code> command is automatically executed at Oracle ACFS startup. When the <code>-o all</code> option is specified, other <code>-o</code> options are ignored. If you want to specify mount options for a registry entry, include those options with the <code>acfsutil registry</code> command when you add the entry to the registry. ▪ <code>ro</code> Mounts the file system in read-only mode. ▪ <code>norootsuid</code> Fails the execution of binaries by non-root users whose permissions allow set user Id execution, and are owned by root. An attempt to run these executables as a non-root user fails with a permission denied error. ▪ <code>rootsuid</code> Allows the execution of binaries by non-root users of set user Id files owned by root. This is the default action.
<code>volume_device</code>	Specifies an Oracle ADVM volume device file that has been formatted by <code>mkfs</code> . <code>device</code> is required but can be a dummy value.
<code>dir</code>	Specifies the directory on which this file system gets mounted. This directory must exist before you issue the mount command. The directory is also known as the file system mount point. The <code>dir</code> directory refers to the root of the file system on the device. <code>dir</code> is required but can be a dummy value.

`mount` is used to attach a file system to the Oracle ACFS hierarchy at the `dir` mount point that is the path name of a directory. The mount happens on the node where the mount command was issued. The mount command returns an error if the file system is not in a dismounted state on this node.

It is not always possible to return the cause of a mount failure to the `mount` command. When this happens Oracle ACFS writes the cause of the failure to the system console and associated system log file.

After `mount` executes successfully, the `MOUNTPATH` field in the `V$ASM_VOLUME` view displays the directory name on which the file system is now mounted.

The `mount` command lists all mounted file systems if it is invoked with no parameters. `root` privilege is required to run `mount`.

Examples

The first example shows how to mount `volume1-123` on the mount point `/u01/app/acfsmounts/myacfs`. The second example shows how to mount all the

registered Oracle ACFS file systems. The dummy names (`none`) have been entered for the device and directory as they are required, but not used, when the `all` option is specified.

Example 13–3 Using mount

```
# /bin/mount -t acfs /dev/asm/volume1-123 /u01/app/acfsmounts/myacfs

# /bin/mount -t acfs -o all none none
```

umount

Purpose

Dismounts an Oracle ACFS file system.

Syntax and Description

```
umount [-h]
umount [-v] volume_device | dir
umount -a [-t acfs]
```

Table 13–5 contains the options available with the `umount` command.

Table 13–5 Options for the umount command

Option	Description
<code>-h</code>	Displays the usage help text and exits.
<code>-t acfs</code>	Specifies the type of file system on Linux. <code>acfs</code> designates the Oracle ACFS type.
<code>-v</code>	Specifies verbose mode. The progress is displayed as the operation occurs.
<code>-a</code>	Specifies to dismount all Oracle ACFS file systems on this node.
<code>volume_device</code>	Specifies an Oracle ADVM volume device file that has been formatted by <code>mkfs</code> .
<code>dir</code>	Specifies the directory on which this file system gets mounted. This directory must exist before you issue the mount command. The director is also known as the file system mount point. The <code>dir</code> directory refers to the root of the file system on the device.

`umount` detaches an Oracle ACFS from the file system hierarchy on the current node. If the file system is busy, `umount` fails.

`root` privileges are required to run the `umount` command.

Examples

The following examples show how to dismount an Oracle ACFS file system. The first example uses the volume device file and the second example uses the file system.

Example 13–4 Using umount

```
# /bin/umount /dev/asm/volume1-123

# /bin/umount /u01/app/acfsmounts/myacfs
```

Oracle ACFS Command-line Tools for Windows Environments

Table 13–6 contains a summary of the Oracle ACFS file system commands for Windows. These commands have been extended to support Oracle ACFS. All other Windows file system commands operate without change for Oracle ACFS.

Table 13–6 Summary of ACFS File System Commands for Windows

Command	Description
<code>acfschkdsk</code>	Checks and repairs an Oracle ACFS file system.
<code>acfsdismount</code>	Dismounts an Oracle ACFS and removes its name space attachment.
<code>acfsformat</code>	Creates an Oracle ACFS file system.
<code>acfsmountvol</code>	Mounts an Oracle ACFS file system.
<code>acfsutil detach</code>	Prepares the Oracle ACFS driver for module unload.

Table 13–7 contains a summary of the non-Oracle ACFS file system commands for Windows.

Table 13–7 Summary of Non-ACFS File System Commands for Windows

Command	Description
<code>advmutil dismount</code>	Dismounts Oracle ASM volume devices for file systems that are not Oracle ACFS.
<code>advmutil list</code>	Lists Oracle ASM volume devices for file systems that are not Oracle ACFS.
<code>advmutil mount</code>	Mounts Oracle ASM volume devices for file systems that are not Oracle ACFS.

acfschkdsk

Purpose

Checks and repairs an Oracle ACFS file system.

Syntax and Description

```
acfschkdsk [/h]
acfschkdsk [/a|/f] [/v] [/n|/y] volume_device
```

Table 13–8 contains the options available with the `acfschkdsk` command.

Table 13–8 Options for the acfschkdsk command

Option	Description
<code>/a</code>	Specifies to automatically fix the file system.
<code>/f</code>	Forces the file system into mountable state without completing a file system check or fix.
<code>/v</code>	Specifies verbose mode. The progress is displayed as the operation occurs.
<code>/h</code>	Displays the usage help text and exits.
<code>/n</code>	Answers no to any prompts.
<code>/y</code>	Answers yes to any prompts.

Table 13–8 (Cont.) Options for the `acfschkdsk` command

Option	Description
<code>volume_device</code>	Specifies the Oracle ADVM volume device.

`acfschkdsk` checks and repairs an existing Oracle ACFS. This command can only be run on a file system that has been dismounted clusterwide.

The Oracle ACFS driver must be loaded for `acfschkdsk` to work. If the driver is not loaded, the administrator is prompted to ensure this is intentional. For information about loading drivers, see "[Oracle ACFS Drivers Resource Management](#)" on page B-3.

The Oracle ACFS driver normally ensures that `acfschkdsk` is the only user of the file system clusterwide. In extreme cases it may be necessary to fix a file system without loading the driver if the file system automount causes the system to crash. The verifications that are normally done by the driver for exclusive access are bypassed in this case.

By default `acfschkdsk` only checks for and only reports any errors. The `/a` flag must be specified to instruct `acfschkdsk` to fix errors in the file system.

In a few cases, `acfschkdsk` prompts for questions before proceeding to check a file system. For example, if `acfschkdsk` detects that another `acfschkdsk` is in progress on the file system, or if `acfschkdsk` detects that the Oracle ACFS driver is not loaded, or if the file system does not appear to be Oracle ACFS. In checking mode, `acfschkdsk` also prompts if there are transaction logs that have not been processed completely due to an unclean shutdown. To run in a non-interactive mode, include either the `/y` or `/n` options to answer yes or no to any questions.

`acfschkdsk` creates working files before it checks a file system. These working files are created in the `temp` directory if space is available. If insufficient space is available `acfschkdsk` attempts to write in the current working directory. The files `acfschkdsk` creates are roughly the size of the file system being checked divided by 32K. At most two files are allocated. These files are deleted when `acfschkdsk` has finished.

In the event that `acfschkdsk` finds a file or directory in the file system for which it cannot determine its name or intended location (possibly due to a corruption in its parent directory), it places this object in the `\lost+found` directory when `acfschkdsk` is run in fix mode. For security reasons only the Windows Administrator is able to read files in `\lost+found`. If the administrator can later determine the original name and location of the file based on its contents, the file can be moved or copied into its intended location.

The file names in the `\lost+found` directory are in the following formats:

```
parent.id.file.id.time-in-sec-since-1970
parent.id.dir.id.time-in-sec-since-1970
```

The `id` fields are the internal Oracle ACFS numeric identifiers for each file and directory in the file system.

You can use `acfsutil info id id mount_point` to attempt to determine the directory associated with `parent.id`. This directory is assumed to be where the deleted object originated. For information about `acfsutil info`, see "[acfsutil info](#)" on page 13-25.

If the parent directory is not known, the parent `id` field is set to UNKNOWN.

Note: It is not possible to see the contents of the `\lost+found` directory from a snapshot.

Examples

The following example shows how to check an Oracle ACFS file system on Windows.

Example 13–5 Using `acfschkdsk`

```
C:\> acfschkdsk asm-volume1-123
```

acfsdismount

Purpose

Dismounts an Oracle ACFS file system and removes its mount point path (name space attachment).

Syntax and Description

```
acfsdismount [/h]
acfsdismount [/v] {drive_letter | path}
acfsdismount [/v] /all
```

Table 13–9 contains the options available with the `acfsdismount` command.

Table 13–9 Options for the `acfsdismount` command

Option	Description
<code>/v</code>	Specifies verbose mode. The progress is displayed as the operation occurs.
<code>/h</code>	Displays help text and exits.
<code>/all</code>	Specifies to dismount all Oracle ACFSs on this node.
<code>drive_letter</code>	Specifies the drive letter followed by a colon.
<code>path</code>	Specifies the full path of the mount point including the drive letter.

`acfsdismount` removes the mount point path (name space attachment) for the specified file system on the current node and dismounts the file system if it is mounted. When the file system is in the dismounted state, the data is preserved on the device and the file system can be re-mounted with the `acfsmountvol` command.

Windows Administrator privileges are required to use the `acfsdismount` command.

Examples

The following examples show how to dismount an Oracle ACFS file system. The first example dismounts a file system using only the drive letter. The second example dismounts a file system using the full path and enables the verbose mode.

Example 13–6 Using `acfsdismount`

```
C:\> acfsdismount O:
```

```
C:\> acfsdismount /v O:\mnt\
```

acfsformat

Purpose

Creates an Oracle ACFS file system.

Syntax and Description

```
acfsformat [/h]
```

```
acfsformat [/v] [/f] [/b blocksize] [/n name] volume_device [blocks]
```

Table 13–10 contains the options available with the `acfsformat` command.

Table 13–10 Options for the `acfsformat` command

Option	Description
<code>/v</code>	Specifies verbose mode. The progress is displayed as the operation occurs.
<code>/n <i>name</i></code>	Specifies the name for the file system. This is also known as a volume label. <code>acfsutil info fs</code> returns the name if one was specified. This can also be displayed from the <code>VOL_LABEL</code> column in the <code>V\$ASM_ACFSVOLUMES</code> view.
<code>/f</code>	Specifies the force option. This action creates the file system even if there is an existing Oracle ACFS on the device, although only if the file system is dismounted. This option overwrites structures on the original file system and should be used with caution.
<code>/h</code>	Displays the usage help text.
<code>/b <i>blocksize</i></code>	The default block size is 4K and this is the only size supported in 11g Release 2 (11.2).
<code><i>blocks</i></code>	Specifies the number of blocks that the file system should consume on the named device. The quantity specified can be in units of K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes). If a unit is not specified, the default is bytes. If number of blocks specified is not a multiple of the block size, than the value is rounded up to the closest multiple. If this option is not specified, the entire device is consumed.
<code><i>volume_device</i></code>	Specifies an Oracle ADVM device file to be formatted.

`acfsformat` creates the on-disk structure needed for Oracle ACFS to be mounted. `acfsformat` performs a quick format on the target volume. After `acfsformat` executes successfully, the `USAGE` column in the `V$ASM_VOLUME` view displays ACFS.

Windows Administrator privileges are not required to use this command. The minimum file system size is 200 MB. The Oracle ACFS driver must be loaded for `acfsformat` to work.

Examples

This example shows how to create an Oracle ACFS file system.

Example 13–7 Using `acfsformat`

```
C:\> acfsformat asm-volume1-123
```

acfsmountvol

Purpose

Mounts an Oracle ACFS file system.

Syntax and Description

```
acfsmountvol /h
acfsmountvol [/v] /all
acfsmountvol [/v] { drive_letter | path } volume_device
```

Table 13–11 contains the options available with the `acfsmountvol` command.

Table 13–11 Options for the `acfsmountvol` command

Option	Description
<code>/v</code>	Specifies verbose mode. The progress is displayed as the operation occurs.
<code>/h</code>	Displays the help usage text and exits.
<code>/all</code>	Reads the ACFS mount registry created with <code>acfsutil registry</code> and mounts the file systems listed in it on this node.
<i>drive_letter</i>	The Windows drive letter to link to an Oracle ADVM volume device.
<i>path</i>	Specifies the full path of the mount point including the drive letter.
<i>volume_device</i>	Specifies an Oracle ADVM volume device file that has been formatted by <code>acfsformat</code> .

`acfsmountvol` is used to attach an Oracle ACFS to the file system hierarchy at the specified *path* or drive letter. *path* must be an empty directory. Oracle ACFS mount points can be created on any empty directory and they can be hierarchical (nested).

After `acfsmountvol` executes successfully, the `MOUNTPATH` column in the `V$ASM_VOLUME` view displays the path or drive letter on which the file system is now mounted. If no arguments are specified, the command lists all mounted Oracle ACFSs on this node.

It is not always possible to return the cause of a mount failure to the `acfsmountvol` command. When this happens Oracle ACFS writes the cause of the failure to the Windows system event logger.

Windows Administrator privileges are required to mount an Oracle ACFS.

Examples

The first example shows how to mount `asm-volume1-123` on the mount point identified as the `F:` drive. The second example shows how to mount `asm-volume1-215` on the mount point `G:\mnt\` and enables verbose mode.

Example 13–8 Using `acfsmountvol`

```
C:\> acfsmountvol F: asm-volume1-123

C:\> acfsmountvol /v G:\mnt\ asm-volume1-215
```

acfsutil detach

Purpose

Prepares the Oracle ACFS driver for module unload.

Syntax and Description

```
acfsutil detach
```

`acfsutil detach` prepares Oracle ACFS for having its driver unloaded. Oracle ACFS registers itself as a base file system upon driver load. In order for the Windows service control manager to issue a stop command to the driver, Oracle ACFS must unregister itself as a base file system in addition to deleting any device objects associated with the driver.

After `acfsutil detach` has completed successfully, and all volumes have been dismounted, then the administrator can issue the Windows service control manager's stop command to the Oracle ACFS driver. If the `acfsutil detach` command is not used, then the stop command does not call the Oracle ACFS driver unload routine. This command can only be invoked by the Windows Administrator. It is called automatically by the Oracle ACFS shutdown scripts.

Examples

The following examples show how to detach all Oracle ACFS file systems and unload drivers on Windows. The first example dismounts all the Oracle ACFS file systems volume devices. The second example detaches the volume devices.

Example 13–9 Using `acfsutil detach`

```
C:\> acfsdismount /all
```

```
C:\> acfsutil detach
```

advmutil dismount

Purpose

The `advmutil dismount` command dismounts Oracle ADVM volume devices for file systems that are not Oracle ACFS.

Syntax and Description

```
advmutil dismount [/unregister] {drive_letter | path}
advmutil dismount [/unregister] [/all | /dg=disk_group]
```

Table 13–12 contains the options available with the Windows `advmutil dismount` command.

Table 13–12 Options for the `advmutil dismount` command

Option	Description
<i>drive_letter</i>	Specifies the drive letter to dismount.
<i>path</i>	Specifies the full path of the mount point, including the drive letter, to dismount.
/all	Specifies to dismount all file systems on this node.

Table 13–12 (Cont.) Options for the `advmutil dismount` command

Option	Description
<code>/unregister</code>	Optional unregister from the system registry an Oracle ADVM volume device and its drive letter. Upon startup of the Oracle ASM instance and Oracle ACFS, the unregistered drive letter and Oracle ADVM volume symbolic links are not created.
<code>/dg=disk_group</code>	Optional name of the disk group containing the registered Oracle ADVM volume devices to create drive letters or possibly mount or dismount file systems.

Note: To format an Oracle ADVM volume device with NTFS, it must be mounted on a drive letter. After formatting, Oracle ADVM volume devices with NTFS can be mounted on either a drive letter or a directory.

To successfully execute this command, the local Oracle ASM instance must be running and the disk groups required by this command must have been created and mounted in the Oracle ASM instance and volumes enabled.

Examples

These examples show how to dismount a file system that is not Oracle ACFS.

1. The first example shows how to remove a drive letter for an Oracle ADVM volume device which disallows access to the file system.
2. The second example shows how to remove a drive letter for an Oracle ADVM volume device which disallows access to the file system and removes it from the registry.
3. The third example shows how to remove a drive letter for all registered Oracle ADVM volume devices which disallows access to all file systems.
4. The fourth example shows how to remove a drive letter for all registered Oracle ADVM volume devices in a particular Oracle ASM disk group which disallows access to all related file systems.

Example 13–10 Using `advmutil dismount`

```
C:\> advmutil dismount H:

C:\> advmutil dismount H: /unregister

C:\> advmutil dismount /all

C:\> advmutil dismount /dg=data
```

advmutil list

Purpose

The `advmutil list` command lists the Oracle ADVM volume devices that were registered with `advmutil mount /register`.

Syntax and Description

```
advmutil list /registry
```

Table 13–13 contains the options available with the Windows `advmutil list` command.

Table 13–13 Options for the `advmutil list` command

Option	Description
<code>list /registry</code>	Lists the drive letter and corresponding Oracle ADVM volume device name contained in the Windows registry.

To successfully execute this command, the local Oracle ASM instance must be running and the disk groups required by this command must have been created and mounted in the Oracle ASM instance and volumes enabled.

For information about listing Oracle ADVM volume devices for Oracle ACFS, see "[acfsutil info](#)" on page 13-25.

Examples

The following example shows how to list all entries in the system registry.

Example 13–11 Using `advmutil list`

```
C:\> advmutil list /registry

Drive Letter      Asm Volume
=====
H:                asm-volume1-123
K:                asm-volume2-245
```

advmutil mount

Purpose

The `advmutil mount` command mounts Oracle ADVM volume devices for file systems that are not Oracle ACFS.

Syntax and Description

```
advmutil mount [/register] {drive_letter | path} volume_device
advmutil mount [/all | /dg=disk_group]
```

Table 13–14 contains the options available with the Windows `advmutil mount` command.

Table 13–14 Options for the `advmutil mount` command

Option	Description
<i>drive_letter</i>	Specifies the drive letter to mount.
<i>path</i>	Specifies the full path of the mount point, including the drive letter, to mount.
<i>volume_device</i>	Indicates the Oracle ADVM volume device.
<code>/all</code>	Specifies to mount all non-Oracle ACFS file systems on this node.

Table 13–14 (Cont.) Options for the `advmutil mount` command

Option	Description
<code>/register</code>	Optional register of an Oracle ADVM volume device and its drive letter in the Windows system registry. After startup of the Oracle ASM instance and Oracle ACFS, all registered drive letter and Oracle ADVM volume symbolic links are created. This allows for automatic file system mounting by Windows as is also done for Oracle ACFS file systems.
<code>/dg=<i>disk_group</i></code>	Optional name of the disk group containing the registered Oracle ADVM volume devices to create drive letters or possibly mount or dismount file systems.

Note: To format an Oracle ADVM volume device with NTFS, it must be mounted on a drive letter. After formatting, Oracle ADVM volume devices with NTFS can be mounted on either a drive letter or a directory.

To successfully execute this command, the local Oracle ASM instance must be running and the disk groups required by this command must have been created and mounted in the Oracle ASM instance and volumes enabled.

For information about registering Oracle ADVM volume devices for Oracle ACFS, see "[acfsutil registry](#)" on page 13-27.

Examples

These examples show how to mount a file system that is not Oracle ACFS.

1. The first example shows how to create a drive letter for an Oracle ASM volume device to be used for non-Oracle ACFS access which can then be used to format the volume possibly for use with NTFS.
2. The second example shows how to create a drive letter for an Oracle ASM volume device to be used for non-Oracle ACFS access which can then be used to format the volume possibly for use with NTFS and save it in the Windows system registry.
3. The third example shows how to create a drive letter for all Oracle ASM volume devices which were saved in the system registry which also allows a file system to be accessible.
4. The fourth example shows how to create a drive letter for all Oracle ASM volume devices in a particular Oracle ASM disk group which were saved in the system registry which also allows a file system to be accessible.

Example 13–12 Using `advmutil mount`

```
C:\> advmutil mount H: asm-volume1-123

C:\> advmutil mount H: asm-volume1-123 /register

C:\> advmutil mount /all

C:\> advmutil mount /dg=data
```

Oracle ACFS Command-Line Tools for All Environments

Table 13–15 contains a summary of the Oracle ACFS commands for all environments. When the options are entered with commands on a Windows platform, use / instead of - with the option.

Table 13–15 Summary of ACFS Commands for All Environments

Command	Description
<code>acfsdbg</code>	Debugs an Oracle ACFS file system.
<code>acfsutil info</code>	Displays various Oracle ACFS file system information.
<code>acfsutil registry</code>	Registers an Oracle ACFS file system with the Oracle ACFS mount registry.
<code>acfsutil rmfs</code>	Removes an Oracle ACFS file system.
<code>acfsutil size</code>	Resizes an Oracle ACFS file system.
<code>acfsutil snap create</code>	Creates a read-only snapshot of an Oracle ACFS file system.
<code>acfsutil snap delete</code>	Deletes a read-only snapshot of an Oracle ACFS file system.
<code>acfsutil tune</code>	Modifies or displays Oracle ACFS tuneables.
<code>advmutil tune</code>	Modifies or displays Oracle ADVN parameters.
<code>advmutil volinfo</code>	Displays information about Oracle ADVN volumes.

acfsdbg

Purpose

Debugs an Oracle ACFS file system.

Syntax and Description

```
acfsdbg [-h]
acfsdbg [-r] [-l] volume_device
```

Table 13–16 contains the options available with the `acfsdbg` command.

Table 13–16 Options for the `acfsdbg` command

Option	Description
<code>-h</code>	Prints out the usage message which displays the various options that can be used when invoking the <code>acfsdbg</code> command, then exits.
<code>-r</code>	Operates in read-only mode. No data is modified on the file system and all write commands are disabled. If the device is mounted anywhere, <code>acfsdbg</code> may not display the latest data as some of it is cached by the file system mounts.
<code>-l</code>	Processes kernel log files. The default is to not process the log files.
<i>volume_device</i>	Specifies the device name of the volume.

`acfsdbg` is an extension to `fsck` and `acfschkdsk`, and is an interactive program that can be used to view and modify on disk structures of the file system. This command should be used with caution by experienced development and support engineers to examine on disk structures to diagnose problems.

When `acfsdbg` is started, it displays a command prompt. At the command prompt, you can enter the subcommands listed in [Table 13–17](#).

`acfsdbg` can also be used by scripts by echoing `acfsdbg` subcommands with a shell pipe to the `acfsdbg` binary.

By default the file system is not modified when running the tool. If the `-l` option is used, the file system metadata in the transaction logs is applied before the interactive disk block dump session. The volume device must specify a volume with a dismounted Oracle ACFS file system. If the volume device has a mounted Oracle ACFS file system, `acfsdbg` displays an error message and exits.

You must be the administrator or a member of the Oracle ASM administrator group to run `acfsdbg`.

Subcommands

[Table 13–17](#) lists the subcommands of `acfsdbg`.

Table 13–17 Subcommands for `acfsdbg`

Option	Description	Syntax
<code>calculate</code>	Calculates simple 2+2 arithmetic expressions Valid operators: + - * / % & ^ ~ << >> White space starts a new expression 0-1 represents a negative 1	<code>calculate [-v] expr [...]</code> -v Verbose mode expr Simple 2+2 expression
<code>cksum</code>	Generates and replaces checksum in header Header offset can be an expression as used by the <code>calculate</code> subcommand White space starts a new header offset Command is disabled in read-only mode	<code>cksum [-c -CE] header_offset [...]</code> -C Regenerate for normal structure checksum -CE Re-generate for Extent structure checksum <i>header_offset</i> Offset of on-disk structure header. The value can be an expression as used by the <code>calculate</code> subcommand
<code>close</code>	Closes the open handle to the device	<code>close</code>
<code>echo</code>	Echoes text on command line to stdout	<code>echo</code>
<code>fenum</code>	Displays the specified File Entry Table (FETA) entry	<code>fenum [-f -e -d] FETA_entry_number</code> -f Displays all on-disk structures related to this structure. -e Displays all on-disk extent information related to this structure. -d Casts the structure as a directory and displays its contents. <i>FETA_entry_number</i> The File Entry Table number used to identify a file on the file system.
<code>help</code>	Displays help message	<code>help</code>
<code>offset</code>	Displays structure at disk offset	<code>offset [-c cast] [-f -d] disk_offset</code> -f Displays all on-disk structures related to this structure. -d Casts the structure as a directory and displays its contents. <i>disk_offset</i> Disk offset to display. The value can be an expression as used by the <code>calculate</code> subcommand.
<code>open</code>	Opens a handle to a device. The default is the volume device name entered on the command line	<code>open [volume_device]</code>

Table 13–17 (Cont.) Subcommands for acfsdbg

Option	Description	Syntax
primary	Sets the context of commands to the primary file system	primary
prompt	Sets the prompt to the specified string	prompt " <i>prompt_string</i> "
quit	Exits the acfsdbg debugger command	quit
read	Reads value from offset The default size to read in is 8 bytes The default count to read is 1	read [-1 -2 -4 -8 -s] [<i>count</i>] <i>offset</i> -1 Read byte value -2 Read 2 byte (short) value -4 Read 4 byte (int) value -8 Read 8 byte (long long) value -s Read null-terminated string <i>count</i> Number of values to read. If not specified, the default is 1 <i>offset</i> Disk offset to read. The value can be an expression as used by the calculate subcommand
snapshot	Sets the context of commands to the specified snapshot	snapshot <i>snapshot_name</i>
write	Writes hexadecimal, octal, or decimal values at the disk offset, estimating how many bytes to write based on value size or number of digits in leading 0 hexadecimal values The disk offset can be an expression used by the calculate subcommand Numeric values can also be an expression as used by the calculate subcommand This command is disabled in read-only mode	write [-1 -2 -4 -8 -c -s] [-C -CE] <i>offset value</i> -1 Write byte value -2 Write 2 byte (short) value -4 Write 4 byte (int) value -8 Write 8 byte (long long) value -c Write text (no null termination). Enclose string in single-quotes (') -s Write null-terminated string. Enclose string in quotes (") -C Regenerate normal structure checksum -CE Regenerate extent structure checksum <i>offset</i> Disk offset to write. The value can be an expression used by the calculate subcommand <i>value</i> The value to write. If numeric, the value can be an expression as used by the calculate subcommand

Examples

These examples show the use of the acfsdbg subcommands.

Example 13–13 Using acfsdbg

```
$ /sbin/acfsdbg /dev/asm/volume1-123
acfsdbg: version = 11.2.0.0.2.0
Oracle ASM Cluster File System (ACFS) On-Disk Structure Version: 38.0
The ACFS volume was created at Mon Aug 2 14:57:45 2009
acfsdbg>

acfsdbg> calculate 60*1024
61,440
61440
61440
0xF000
0170000
1111:0000:0000:0000

acfsdbg> prompt "acfsdbg test>"
```

```
acfsdbg test>
echo "offset 64*1024" | acfsdbg /dev/asm/volume1-123
```

acfsutil info

Purpose

Displays various Oracle ACFS file system information.

Syntax and Description

```
acfsutil info fs [-h ]
acfsutil info fs [-o item] [path...]
acfsutil info file path
acfsutil info id num path
```

[Table 13–18](#) contains the options available with the `acfsutil info` command.

`acfsutil info -h` displays help text and exits.

Table 13–18 Options for the `acfsutil info` command

Option	Description
<code>-o</code>	<p>Displays the specific file system item from the following list:</p> <ul style="list-style-type: none"> ■ <code>freespace</code> - Amount of free space available in the file system ■ <code>totalspace</code> - Size of the file system ■ <code>mountpoints</code> - Mount point associated with <code>path</code> ■ <code>volumes</code> - Number of volumes associated with the file system ■ <code>primaryvolume</code> - Path name of the device associated with file system ■ <code>available</code> - Length of time the file system has been available ■ <code>ismountpoint</code> - 1 if <code>path</code> is the mount point, 0 if <code>path</code> is not a mount point ■ <code>isavailable</code> - 1 if the file system is available, 0 if file system is not available ■ <code>iscorrupt</code> - 1 if the file system is corrupt, 0 if file system is not corrupt ■ <code>diskgroup</code> - Name of the disk group associated with <code>path</code> ■ <code>redundancy</code> - Type of redundancy used by the Oracle ADVM volume ■ <code>resizeincrement</code> - Allocation unit size in megabytes of the Oracle ADVM volume used by <code>path</code> ■ <code>stripewidth</code> - Stripe width for the Oracle ADVM volume used by <code>path</code> ■ <code>stripecolumns</code> - Stripe set columns for the Oracle ADVM volume used by <code>path</code>
<code>path</code>	Specifies an existing directory folder where the mount point resides or a drive letter.

Table 13–18 (Cont.) Options for the `acfsutil info` command

Option	Description
<code>num</code>	Specifies the Oracle ACFS file identifier number reported by the Oracle ACFS driver. The number should be specified in decimal format.

`acfsutil info fs` prints out basic information about Oracle ACFS such as the driver version number and detailed information about the specified mounted file system.

The file system information includes the volume device name, the size of the file system, the amount of space available on the volume device, the file system mount time on this node, the state of the file system, the user specified block size, the number of ACFS snapshots, the space consumed by snapshots in the file system, and the optional name or volume label associated with the file system. Any user can run `acfsutil info fs`.

Possible Oracle ACFS states include:

- `offline` indicates that the underlying devices are not accessible, possibly due to an Oracle ASM instance failure, disk group forced dismount, or an irrecoverable I/O error. The file system on this node can only be dismounted. All other attempts at access result in errors.
- `available` indicates that the file system is on line and operational
- `corrupt` indicates that `fsck` or `acfschkdsk` should be run on the file system at the earliest possible convenience to correct a detected inconsistency

`acfsutil info file` displays information about an Oracle ACFS file such as its extent map which details the locations of the blocks comprising the file. It also displays information such as the used versus allocated storage for a file. You must have read access to the specified file to run `acfsutil info file`.

`acfsutil info id` is used to translate an internal numeric Oracle ACFS file identifier to a path name in the file system. This is useful when the Oracle ACFS driver reports I/O errors to the system event logger associated with a particular file in an Oracle ACFS and identifies it by its internal identifier. You must have administrator privileges or you must be a member of the Oracle ASM administrator group to run `acfsutil info id`.

Examples

The following are examples of the use of `acfsutil info`. The first example displays information about the specified file system. The second example checks whether the file system is corrupt.

Example 13–14 Using `acfsutil info`

```
$ /sbin/acfsutil info fs /u01/app/acfsmounts/myacfs
/u01/app/acfsmounts/myacfs
ACFS Version: 11.2.0.1.0.0
flags:      MountPoint,Available
mount time: Tue Jun 29 10:18:27 2009
volumes:   1
total size: 268435456
total free: 163360768
primary volume: /dev/asm/volume1-228
label:
flags:      Primary,Available,ADVM
on-disk version: 39.0
```

```

allocation unit:      4096
major, minor:        252, 116737
size:                268435456
free:                163360768
ADVM diskgroup       DATA
ADVM resize increment: 268435456
ADVM redundancy:     mirror
ADVM stripe columns: 4
ADVM stripe width:   131072
number of snapshots: 3
snapshot space usage: 106496

$ /sbin/acfsutil info fs -o iscorrupt /u01/app/acfsmounts/myacfs
0

```

acfsutil registry

Purpose

Registers an Oracle ACFS file system with the Oracle ACFS mount registry.

Syntax and Description

```

acfsutil registry [-h]
acfsutil registry
acfsutil registry -a [-f] [-n { nodes | all } ] [-o moptions] device path
acfsutil registry -d {device | path}
acfsutil registry -l [device | path]
acfsutil registry -m device

```

Table 13–19 contains the options available with the `acfsutil registry` command.

`acfsutil registry -h` displays help text and exits.

Table 13–19 Options for the `acfsutil registry` command

Option	Description
-a	Specifies to add the device, path, and associated <i>moptions</i> to the Oracle ACFS mount registry. The Oracle ASM volume device specified must exist on the local node to add the information to the mount registry. The arguments represent all the information needed to mount the file system. At Oracle ACFS startup time these file systems are automatically mounted. On Windows the volume mount points on the specified directories are created if they do not exist. Duplicate device entries are not allowed. Duplicate mount points are allowed but must be used with the <code>-n</code> option for disjoint node-specific mounts.
-d	Deletes the device or path from the Oracle ACFS mount registry. If a path is specified and it is not unique in the Oracle ACFS mount registry, the command fails and you must reissue the command specifying the device as an argument instead.
-f	This is used in combination with <code>-a</code> when the specified device might exist in the registry and the administrator wants to replace the registration.

Table 13–19 (Cont.) Options for the `acfsutil registry` command

Option	Description
<code>-n</code>	This is used in combination with <code>-a</code> and specifies an optional comma-delimited list of nodes, or the <code>all</code> keyword. This specifies which nodes should attempt to mount this device on the specified mount point. Host names should be specified. The <code>all</code> keyword is functionally equivalent to not specifying a list of nodes and indicates that the device should be mounted on all nodes. Mounting on all nodes is the default behavior.
<code>-o options</code>	Specifies the mount options to be used when mounting the file system. For Linux only. All options are available except the <code>all</code> option. See " mount " on page 13-10.
<code>-l</code>	Lists all the mount points currently in the registry on a single line, with fields separated by a colon (:).
<code>-m</code>	List the registered mount point, if one exists, associated with the specified device.
<code>path</code>	Specifies an existing directory folder where the mount point resides or a drive letter.
<code>device</code>	Specifies an Oracle ACFS device file that has been formatted.

`acfsutil registry` adds or deletes a file system from the Oracle ACFS persistent mount registry. The mount registry is a global registry that is used at Oracle ACFS startup on each node to mount all file systems specified in it. `root` or `asmadmin` privileges are required to modify the registry. For information about operating system group privileges, see "[About Privileges for Oracle ASM](#)" on page 3-23.

Any user is allowed to display the contents of the registry. To mount all the file systems in the Oracle ACFS mount registry, use the platform specific mount command with the `all` option. This is done automatically at Oracle ACFS startup on each node.

If no options are specified, the command displays all of the Oracle ACFS mounts in the registry.

See "[About the Oracle ACFS Mount Registry](#)" on page 5-6.

Examples

The following examples show the use of `acfsutil registry`. The first example shows how to add the volume device file and file system mount point to the registry. The second example shows how to list the registered mount point associated with the specified volume device file. The third example shows how to delete the specified volume device file from the registry.

Example 13–15 Using `acfsutil registry`

```
$ /sbin/acfsutil registry -a /dev/asm/volume1-123 /u01/app/acfsmounts/myacfs
$ /sbin/acfsutil registry -m /dev/asm/volume1-123
$ /sbin/acfsutil registry -d /dev/asm/volume1-123
```

acfsutil rmfs

Purpose

Removes an Oracle ACFS file system.

Syntax and Description

```
acfsutil [-h]
acfsutil rmfs device
```

Table 13–20 contains the options available with the `acfsutil rmfs` command.

`acfsutil rmfs -h` displays help text and exits.

Table 13–20 Options for the `acfsutil rmfs` command

Option	Description
<i>device</i>	Specifies an Oracle ACFS device file that has been formatted.

`acfsutil rmfs` is used to remove an Oracle ACFS that is dismounted. When the command is executed, the superblock of the file system is disabled. `root` or `asmadmin` privileges are required to run this command. For information about operating system group privileges, see "About Privileges for Oracle ASM" on page 3-23.

After `acfsutil rmfs` executes successfully, the `MOUNTPATH` and `USAGE` columns in the `V$ASM_VOLUME` view are cleared for the device. The removed Oracle ACFS can be restored using `fsck` or `acfschkdsk`. The device can be reformatted with a new Oracle ACFS using the `mkfs` or `acfsformat` commands.

Examples

The following example shows the use of `acfsutil rmfs` to remove the specified volume device file and associated file system.

Example 13–16 Using `acfsutil rmfs`

```
$ /sbin/acfsutil rmfs /dev/asm/volume1-123
```

acfsutil size**Purpose**

Resizes an Oracle ACFS file system.

Syntax and Description

```
acfsutil size [-h]
acfsutil size [+|-]n[K|M|G|T|P] [device] path
```

Table 13–21 contains the options available with the `acfsutil size` command.

`acfsutil size -h` displays help text and exits.

Table 13–21 Options for the `acfsutil size` command

Option	Description
<code>[+ -]<i>n</i></code>	Specifies the new size for the Oracle ACFS file system where <i>n</i> is a valid positive whole number greater than zero. The number can be preceded by a + or - to indicate the amount to add or decrease. If no operand exists, the new size is the absolute size.
<code>K M G T P</code>	Specifies that the integer supplied for size is in the units of K (Kilobytes), M (Megabytes), G (Gigabytes), T (Terabytes), or P (Petabytes). If the unit indicator is specified, then it must be appended to the integer. If omitted, the default unit is bytes.

Table 13–21 (Cont.) Options for the `acfsutil size` command

Option	Description
<code>device</code>	Specifies the optional volume device file.
<code>path</code>	Specifies an existing directory folder where the mount point resides or a drive letter.

`acfsutil size` grows or shrinks the mounted Oracle ACFS and its underlying Oracle ADVM storage to match the new size specified. This operation also resizes the underlying Oracle ASM Volume file to match the new length that is specified. The disk group must have enough free storage to accommodate any requested increase to the file system size.

Reducing a file system size returns unused storage space located at the end of the file system to the disk group. Shrinking in this release is intended for accidents when the wrong initial size or resize increment was specified, and before the storage was actually used. After storage has been used for user data or file system metadata, it may not be possible to shrink the file system, even if the files using that storage have been deleted.

Note: Defragmenting a file system is not supported in this release; only unused storage can be deallocated from the end of the file system.

The size value is rounded up based on the block size of the file system and the allocation unit of the Oracle ASM volume device file. To determine the Oracle ASM volume device resize increment, examine the `RESIZE_UNIT_MB` field in the `V$ASM_VOLUME` view, or look for `Resize Unit` in the output of `asmcmd volinfo`.

There is a limit of 5 extents for the file system's internal storage bitmap. This cause any attempts to increase the file system to fail after it has been increased four or more times. When the limit on a file system expansion has been reached, running `fsck` or `acfschkdsk` with the `-a` option may consolidate the internal storage bitmap, allowing future file system expansion.

`root` or users who are members of the `asmadmin` group can issue this command. For information about operating system group privileges, see "[About Privileges for Oracle ASM](#)" on page 3-23.

Examples

The following example shows the use of `acfsutil size`. This example increases the primary device file of `/u01/app/acfsmounts/myacfs` file system by 500 MB.

Example 13–17 Using `acfsutil size`

```
$ /sbin/acfsutil size +500M /u01/app/acfsmounts/myacfs
```

acfsutil snap create

Purpose

Creates a read-only snapshot of an Oracle ACFS file system.

Syntax and Description

```
acfsutil snap create [-h]
```



```
acfsutil snap create snapshot mount_point
```

[Table 13–22](#) contains the options available with the `acfsutil snap create` command.

`acfsutil snap create -h` displays help text and exits.

Table 13–22 Options for the `acfsutil snap create` command

Option	Description
<i>snapshot</i>	Specifies a name for the snapshot. The <code>.ACFS/snaps</code> directory itself cannot be snapped.
<i>mount_point</i>	Specifies the mount point.

`acfsutil snap create` creates a read-only snapshot of the Oracle ACFS mounted on *mount_point*. Snapshots are not separate file systems. The snapshot appears in `.ACFS/snaps/snapshot` and is a complete replica of the file system at the time the snapshot command was given. Snapshots usually use very little storage initially as they share file system blocks with the original file system until a file changes.

Tools such as `du` report the total disk space usage of the snapshotted files, which includes the storage shared with the original versions of the files. To determine the total space used for the snapshots, use the `acfsutil info fs` command. See "[acfsutil info](#)" on page 13-25.

Oracle ACFS snapshots are immediately available for use after they are created. They are always online under the `.ACFS/snaps` directory when the original file system is mounted. No separate command is needed to mount them.

The last modification time of `.ACFS/snaps/snapshot` is set to the time the snapshot was created, not the last modification time of the root directory of the file system at the time of the snapshot creation. An `ls -l` or similar command in the `.ACFS/snaps` directory can be used to determine the age of each snapshot.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

For more information about Oracle ACFS snapshots, see "[About Oracle ACFS Snapshots](#)" on page 5-6.

Examples

The following examples show the use of `acfsutil snap create`. The first and second examples create snapshots of the Oracle ACFS file system mounted on `/u01/app/acfsmounts/myacfs`. The third example lists the contents of the `.ACFS/snaps` directory.

Example 13–18 Using `acfsutil snap create`

```
$ /sbin/acfsutil snap create mysnapshot_0900609b /u01/app/acfsmounts/myacfs

$ /sbin/acfsutil snap create mysnapshot_0900609c /u01/app/acfsmounts/myacfs

$ ls /u01/app/acfsmounts/myacfs/.ACFS/snaps
mysnapshot_0900609a  mysnapshot_0900609b  mysnapshot_0900609c
```

acfsutil snap delete

Purpose

Deletes a read-only snapshot of an Oracle ACFS file system.

Syntax and Description

```
acfsutil snap delete [-h]
acfsutil snap delete snapshot mount_point
```

Table 13–23 contains the options available with the `acfsutil snap delete` command.

`acfsutil snap delete -h` displays help text and exits.

Table 13–23 Options for the `acfsutil snap delete` command

Option	Description
<i>snapshot</i>	Specifies a name for the snapshot.
<i>mount_point</i>	Specifies the mount point of the file system that contains the snapshot.

`acfsutil snap delete` deletes the snapshot named *snapshot* in the Oracle ACFS mounted on *mount_point*. After successful completion of the command, the representation of the snapshot in the `.ACFS/snaps` directory is removed. The command fails if any file within the snapshot is open on any cluster node.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

The following example deletes a snapshot of the file system mounted on `/u01/app/acfsmounts/myacfs`.

Example 13–19 Using `acfsutil snap delete`

```
$ /sbin/acfsutil snap delete mysnapshot_0900609b /u01/app/acfsmounts/myacfs
```

acfsutil tune

Purpose

The `acfsutil tune` command displays the value of a specific tuneable parameter or all Oracle ACFS tuneable parameters, or sets the value of a tuneable in a persistent manner on a particular node.

Syntax and Description

```
acfsutil [-h]
acfsutil tune [tuneable_name]
acfsutil tune tuneable_name=value
```

Table 13–24 contains the options available with the `acfsutil tune` command.

`acfsutil tune -h` displays help text and exits.

Table 13–24 Options for the `acfsutil tune` command

Option	Description
<i>tuneable_name</i>	Specifies the name of the tuneable parameter.
<i>value</i>	Specifies the value for a tuneable parameter.

The only Oracle ACFS tuneable parameter is the Windows specific `AcfsMaxOpenFiles`, which limits the number of open Oracle ACFS files on Windows. Normally you do not need to change the value of this tuneable.

Changing a tuneable parameter has an immediate effect and persists across restarts. You must be a root user or the Windows Administrator to change the value of a tuneable.

Examples

The first example prints Oracle ACFS tuneables. The second example changes the value of a tuneable.

Example 13–20 Using `acfsutil tune`

```
$ /sbin/acfsutil tune
$ acfsutil tune AcfsMaxOpenFiles=50000
```

advmutil tune

Purpose

`advmutil tune` displays the value of a specific Oracle ADVM parameter or sets the value of a specific Oracle ADVM parameter.

Syntax and Description

```
advmutil [-h]
advmutil tune parameter [= value]
```

[Table 13–25](#) contains the options available with the `advmutil tune` command.

`advmutil tune -h` displays help text and exits.

Table 13–25 Options for the `advmutil tune` command

Option	Description
<i>parameter</i>	Specifies the parameter for which you want to set or display the value.
<i>value</i>	Optional value provided to set the value of the specified parameter.

If a value is not provided, the `advmutil tune` command displays the value that is currently assigned to the specified parameter.

The parameters that can be specified with `advmutil tune` are the maximum kernel memory (`max_memory`) or the maximum time in minutes for the deadlock timer (`deadlock_timer`).

The maximum kernel memory (`max_memory`) specifies the maximum operating system (OS) kernel memory in megabytes that can be consumed by the Oracle ASM Dynamic Volume Manager driver to cache Oracle ASM extent maps.

Note: The `deadlock_timer` parameter should be only be set by Oracle Support Services.

Examples

The first example sets the maximum kernel memory. The second example queries the current setting for the maximum kernel memory to be consumed by the Oracle ADVM. The third example changes the maximum time in minutes for the deadlock timer. The fourth example queries the current setting of a parameter.

Example 13–21 Using `advmutil tune`

```
$ /sbin/advmutil tune max_memory = 40

$ /sbin/advmutil tune max_memory
max_memory = 40 (0x28)

$ /sbin/advmutil tune deadlock_timer = 20

$ /sbin/advmutil tune deadlock_timer
deadlock_timer = 20 (0x14)
```

advmutil volinfo

Purpose

`advmutil volinfo` displays information about Oracle ADVM volume devices.

Syntax and Description

```
advmutil volinfo [-h]
advmutil volinfo [-l] volume_device
```

Table 13–26 contains the options available with the `advmutil volinfo` command.

`advmutil volinfo -h` displays help text and exits.

Table 13–26 Options for the `advmutil volinfo` command

Option	Description
<i>volume_device</i>	Specifies the volume device name.
-l	Optional flag separates the Oracle ADVM volume device information by colons.

`advmutil volinfo` displays information about the specified Oracle ADVM volume device in a list format by default. The `-l` option on Linux (`/l` on Windows) formats the display into a colon-separated string.

Examples

The first example displays information about an Oracle ADVM volume device, using the `advmutil volinfo` command with the volume device name. The second

example displays information about the volume device in a format suitable for scripting, using the `-l` option with `advmutil volinfo`.

Example 13–22 Using `advmutil volinfo`

```
$ /sbin/advmutil volinfo /dev/asm/volume1-123
```

```
Interface Version: 1
Size (MB): 256
Resize Increment (MB): 256
Redundancy: mirror
Stripe Columns: 4
Stripe Width (KB): 128
Disk Group: DATA
Volume: VOLUME1
```

```
$ /sbin/advmutil volinfo -l /dev/asm/volume1-228
```

```
Interface Version : 1 : Size (MB) : 256 : Resize Increment (MB) : 256 :
Redundancy : mirror : Stripe Columns : 4 : Stripe Width (KB) : 128 :
Disk Group : DATA : Volume : VOLUME1
```

Creating Database Files in Oracle ASM Using SQL*Plus

This appendix describes how to deploy Oracle Automatic Storage Management (Oracle ASM) using SQL*Plus under the following topics:

- [Creating a Database that Uses Oracle ASM](#)
- [Creating Tablespaces in Oracle ASM](#)
- [Creating Redo Logs in Oracle ASM](#)
- [Creating Control Files in Oracle ASM](#)
- [Creating Archive Log Files in Oracle ASM](#)

See Also: *Oracle Database Administrator's Guide* for information about using Oracle Managed Files (OMF)

Creating a Database that Uses Oracle ASM

The recommended method of creating your database is to use the Database Configuration Assistant (DBCA). However, if you create your database manually using the `CREATE DATABASE` statement, then Oracle ASM enables you to create a database and all of its underlying files with a minimum of input.

The following section describes an example of using the `CREATE DATABASE` statement, where database files are created and managed automatically by Oracle ASM.

Creating a Database that Uses Oracle ASM: Example

This example creates a database with the following Oracle ASM files:

- A `SYSTEM` tablespace data file in disk group `data`.
- A `SYSAUX` tablespace data file in disk group `data`. The tablespace is locally managed with automatic segment-space management.
- A multiplexed online redo log with two online log groups, one member of each in `data` and `fra` (fast recovery area).
- If automatic undo management mode is enabled, then an undo tablespace data file in directory `data`.
- If you do not specify a value for the `CONTROL_FILES` initialization parameter, then two control files, one in `data` and another in `fra` (fast recovery area). The control file in `data` is the primary control file.

The following initialization parameter settings are included in the initialization parameter file:

```
DB_CREATE_FILE_DEST = +data
DB_RECOVERY_FILE_DEST = +fra
DB_RECOVERY_FILE_DEST_SIZE = 10G
```

The following statement is issued at the SQL prompt:

```
CREATE DATABASE orcl;
```

Creating Tablespaces in Oracle ASM

When Oracle ASM creates a data file for a permanent tablespace (or a temporary file for a temporary tablespace), the data file is set to auto-extensible with an unlimited maximum size and 100 MB default size. You can use the `AUTOEXTEND` clause to override this default extensibility and the `SIZE` clause to override the default size.

Oracle ASM applies attributes to the data file, as specified in the system default template for a data file as shown in the table in ["Managing Disk Group Templates"](#) on page 7-15. You can also create and specify your own template.

Files in a tablespace might be in both Oracle ASM files and non-Oracle ASM files as a result of the tablespace history. RMAN commands enable non-Oracle ASM files to be relocated to an Oracle ASM disk group and enable Oracle ASM files to be relocated as non-Oracle ASM files.

For example, if there is a disk group named `data`, you can create a tablespace `tblspace` in that disk group with the following SQL statement:

```
CREATE TABLESPACE tblspace DATAFILE +data;
```

The following example illustrates the usage of Oracle ASM with defaults. This example enables Oracle ASM to create and manage the tablespace data file for you, using Oracle supplied defaults that are adequate for most situations.

Assume the following initialization parameter setting:

```
DB_CREATE_FILE_DEST = +data
```

The following statement creates the tablespace and its data file:

```
CREATE TABLESPACE tblspace;
```

Creating Redo Logs in Oracle ASM

Online redo logs can be created in multiple disk groups, either implicitly in an initialization parameter file or explicitly in an `ALTER DATABASE . . . ADD LOGFILE` statement. Each online log should have one log member in multiple disk groups. The filenames for log file members are automatically generated. All partially created redo log files, created as a result of a system error, are automatically deleted.

For example, if there are disk groups named `data1` and `data2`, you can create a log file with a member in each disk group with the following SQL statement:

```
ALTER DATABASE ADD LOGFILE (+data1,+data2);
```

The following example creates a log file with a member in each of the disk groups `data1` and `data2`. The following parameter settings are included in the initialization parameter file:

```
DB_CREATE_ONLINE_LOG_DEST_1 = +data1
DB_CREATE_ONLINE_LOG_DEST_2 = +data2
```


The following statement is issued at the SQL prompt:

```
ALTER DATABASE ADD LOGFILE;
```

Creating Control Files in Oracle ASM

Control files can be explicitly created in multiple disk groups. The filenames for control files are automatically generated.

If you must specify a control file by name, then you can use alias filenames to reference Oracle ASM files and use easy to understand names. Using aliases in the specification of a control file during its creation enables you to later refer to the control file with a user-specified name. An alias can also be specified as a control file name in the `CONTROL_FILES` initialization parameter.

When creating a control file, data files and log files that are stored in Oracle ASM disk groups should be given to the `CREATE CONTROLFILE` command using the file reference form of their Oracle ASM filenames. However, the use of the `RESETLOGS` option requires the use of a file creation form for the specification of the log files.

Creating Control Files in Oracle ASM: Example 1

The following `CREATE CONTROLFILE` statement is generated by an `ALTER DATABASE BACKUP CONTROLFILE TO TRACE` command for a database with data files and log files created on disk groups `data` and `fra`:

```
CREATE CONTROLFILE REUSE DATABASE "ORCL" NORESETLOGS ARCHIVELOG
  MAXLOGFILES 16
  MAXLOGMEMBERS 3
  MAXDATAFILES 100
  MAXINSTANCES 8
  MAXLOGHISTORY 292
LOGFILE
  GROUP 1 (
    '+DATA/orcl/onlinelog/group_1.261.685366203',
    '+FRA/orcl/onlinelog/group_1.257.685366205'
  ) SIZE 50M BLOCKSIZE 512,
  GROUP 2 (
    '+DATA/orcl/onlinelog/group_2.262.685366207',
    '+FRA/orcl/onlinelog/group_2.258.685366211'
  ) SIZE 50M BLOCKSIZE 512,
  GROUP 3 (
    '+DATA/orcl/onlinelog/group_3.263.685366213',
    '+FRA/orcl/onlinelog/group_3.259.685366215'
  ) SIZE 50M BLOCKSIZE 512
-- STANDBY LOGFILE
DATAFILE
  '+DATA/orcl/datafile/system.256.685366089',
  '+DATA/orcl/datafile/sysaux.257.685366089',
  '+DATA/orcl/datafile/undotbs1.258.685366091',
  '+DATA/orcl/datafile/users.259.685366091'
CHARACTER SET WE8MSWIN1252
;
```

Creating Control Files in Oracle ASM: Example 2

This example is a `CREATE CONTROLFILE` statement for a database with data files, but uses a `RESETLOGS` clause, and thus uses the creation form for log files:

```

CREATE CONTROLFILE REUSE DATABASE "ORCL" RESETLOGS ARCHIVELOG
    MAXLOGFILES 16
    MAXLOGMEMBERS 3
    MAXDATAFILES 100
    MAXINSTANCES 8
    MAXLOGHISTORY 292
LOGFILE
GROUP 1 (
    '+DATA/orcl/onlinelog/group_1.261.685366203',
    '+FRA/orcl/onlinelog/group_1.257.685366205'
) SIZE 50M BLOCKSIZE 512,
GROUP 2 (
    '+DATA/orcl/onlinelog/group_2.262.685366207',
    '+FRA/orcl/onlinelog/group_2.258.685366211'
) SIZE 50M BLOCKSIZE 512,
GROUP 3 (
    '+DATA/orcl/onlinelog/group_3.263.685366213',
    '+FRA/orcl/onlinelog/group_3.259.685366215'
) SIZE 50M BLOCKSIZE 512
-- STANDBY LOGFILE
DATAFILE
    '+DATA/orcl/datafile/system.256.685366089',
    '+DATA/orcl/datafile/sysaux.257.685366089',
    '+DATA/orcl/datafile/undotbs1.258.685366091',
    '+DATA/orcl/datafile/users.259.685366091'
CHARACTER SET WE8MSWIN1252
;

```

Creating Archive Log Files in Oracle ASM

Disk groups can be specified as archive log destinations in the `LOG_ARCHIVE_DEST` and `LOG_ARCHIVE_DEST_n` initialization parameters. When destinations are specified in this manner, the archive log filename will be unique, even if archived twice.

If `LOG_ARCHIVE_DEST` is set to a disk group name, `LOG_ARCHIVE_FORMAT` is ignored. Unique filenames for archived logs are automatically created by the Oracle database. If `LOG_ARCHIVE_DEST` is set to a directory in a disk group, `LOG_ARCHIVE_FORMAT` has its normal semantics.

The following sample archive log names might be generated with `DB_RECOVERY_FILE_DEST` set to `+data`. `ORCL` is the value of the `DB_UNIQUE_NAME` parameter:

```

+DATA/ORCL/ARCHIVELOG/2009_09_23/thread_1_seq_38.614.541956473
+DATA/ORCL/ARCHIVELOG/2009_09_23/thread_4_seq_35.609.541956477
+DATA/ORCL/ARCHIVELOG/2009_09_23/thread_2_seq_34.603.541956487
+DATA/ORCL/ARCHIVELOG/2009_09_25/thread_3_seq_100.621.541956497
+DATA/ORCL/ARCHIVELOG/2009_09_25/thread_1_seq_38.614.541956511

```

Oracle ACFS Advanced Topics

This appendix discusses limits, advanced administration, and troubleshooting for Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

This appendix contains the following topics:

- [Limits of Oracle ACFS](#)
- [Limits of Oracle ADVM](#)
- [Oracle ACFS Drivers Resource Management](#)
- [Oracle ACFS Registry Resource Management](#)
- [Oracle ACFS Individual File System Resource Management](#)
- [Oracle ACFS and Oracle Restart](#)
- [Understanding Oracle ACFS I/O Failure Console Messages](#)

For information about Oracle ACFS, see [Chapter 5, "Introduction to Oracle ACFS"](#).

Limits of Oracle ACFS

The limits of Oracle ACFS and these discussed in this section.

The topics contained in this section are:

- [Oracle ACFS Disk Space Usage](#)
- [Oracle ACFS Error Handling](#)
- [Oracle ACFS and NFS](#)

Oracle ACFS Disk Space Usage

Oracle ACFS supports 64 million files in a file system, 63 snapshots, up to 64 mounts on 32 bit systems, and 256 mounts on 64 bit systems.

Oracle ACFS preallocates large user files to improve performance when writing data. This storage is not returned when the file is closed, but it is returned when the file is deleted. Oracle ACFS also allocates local metadata files as nodes mount the file system for the first time. This can result in a mount failing due to an out of space error, and much of this storage must be contiguous. This storage is approximately 64-128 megabyte per node. Oracle ACFS also keeps local bitmaps available to reduce contention on the global storage bitmap when searching for free space. This disk space is reported as `in use` by tools such as the UNIX `df` command even though some of it may not actually be allocated as of yet. This local storage pool can be as large as 128 megabytes per node.

Oracle ACFS Error Handling

Oracle ASM instance failure or forced shutdown while Oracle ACFS or another file system is using an Oracle ADVM volume results in I/O failures. The volumes must be closed and re-opened to access the volume again. This requires dismounting any file systems that were mounted when the local Oracle ASM instance failed. After the instance is restarted, the corresponding disk group must be mounted with volume enabled followed by a remount of the file system. See "[Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#)" on page 13-5.

If any file systems are currently mounted on Oracle ADVM volume files, the `SHUTDOWN ABORT` command should not be used to terminate the Oracle ASM instance without first dismounting those file systems. Otherwise, applications encounter I/O errors and Oracle ACFS user data and metadata being written at the time of the termination may not be flushed to storage before the Oracle ASM storage is fenced. If there is not time to permit the file system to dismount, then you should issue two `sync (1)` commands to flush cached file system data and metadata to persistent storage before issuing the `SHUTDOWN ABORT` operation.

Oracle ACFS does not interrupt the operating system environment when a metadata write fails, whether due to Oracle ASM instance failure or storage failure. Instead, Oracle ACFS isolates errors to an individual file system, putting it in an offline error state. The only operation that succeeds on that node for that file system from that point forward is a dismount operation. Another node recovers any outstanding metadata transactions, assuming it can write the metadata out to the storage. It is possible to remount the file system on the offlined node after the I/O condition is resolved.

It might not be possible for an administrator to dismount a file system while it is in the offline error state if there are processes referencing the file system, such as a directory of the file system being the current working directory for a process. To dismount the file system in this case it would be necessary to identify all processes on that node with references to files and directories on the file system and cause them to exit. The Linux `fuser` or `lsof` commands or Window `handle` command list information about processes and open files.

If Oracle ACFS detects inconsistent file metadata returned from a read operation, based on checksum or expected type comparisons, Oracle ACFS takes the appropriate action to isolate the affected file system components and generate a notification that `fsck` or `acfschkdsk` should be run as soon as possible. Each time the file system is mounted a notification is generated with a system event logger message until `fsck` or `acfschkdsk` is run.

Oracle ACFS and NFS

When exporting file systems through NFS on Linux, use the `-fsid=num` exports option. This option forces the file system identification portion of the file handle used to communicate with NFS clients to be the specified number instead of a number derived from the major and minor number of the block device on which the file system is mounted. Any 32-bit number can be used for `num` but it must be unique among all the exported file systems. In addition, `num` must be unique among members of the cluster and must be the same `num` on each member of the cluster for a given file system. This is needed because Oracle ASM DVM block device major numbers are not guaranteed to be the same across reboots of the same node or across different nodes in the cluster.

Limits of Oracle ADVM

The limits of Oracle ADVM and these discussed in this section.

The default configuration for an Oracle ADVM volume is four columns of 64 MB extents in length and a 128 KB stripe width. Oracle ADVM writes data as 128 KB stripe chunks in round robin fashion to each column and fills a stripe set of four 64 MB extents with 2000 stripe chunks before moving to a second stripe set of four 64 MB extents for volumes greater than 256 megabytes. Note that setting the number of columns on an Oracle ADVM dynamic volume to 1 effectively turns off striping for the Oracle ADVM volume.

On Linux platforms Oracle ASM Dynamic Volume Manager (Oracle ADVM) volume devices are created as block devices regardless of the configuration of the underlying storage in the Oracle ASM disk group. Do not use `raw (8)` to map Oracle ADVM volume block devices into raw volume devices.

Oracle ACFS Drivers Resource Management

The Oracle ACFS drivers resource is supported only for Oracle grid infrastructure cluster configurations; it is not supported for Oracle Restart configurations. See ["Oracle ACFS and Oracle Restart"](#) on page B-6.

The Oracle ACFS drivers resource (`ora.drivers.acfs`) is created by the Grid Infrastructure root script that is executed following the Grid Infrastructure installation. The Oracle ASM instance resource (`ora.asm`) names the drivers resource as a weak dependency. As a result, the start action for the drivers resource is also called whenever the start action for the `ora.asm` resource is issued. The start action for the drivers resource includes support for loading the Oracle ACFS, Oracle ADVM, and OKS drivers into the operating system.

Following an Oracle grid infrastructure installation on Linux and UNIX platforms, a root script is executed that includes actions for copying the Oracle ACFS components; including the Oracle ACFS, Oracle Kernel Services Driver (OKS), and Oracle ADVM drivers; into operating system-specific locations.

- On Linux systems, the command-line tools are copied into the `/sbin` directory. The drivers are copied into the `extra/usm` directory under the appropriate `/lib/modules/Linux_kernel_version`. For example:

```
$ acfsdriverstate version
ACFS-9205: OS/ADVM,ACFS installed version = 2.6.18-8.el5xen(i386)/090712
```

```
$ ls /lib/modules/2.6.18-8.el5xen/extra/usm
oracleacfs.ko oracleadv.ko oracleoks.ko
```

- On Windows, the command-line tools are copied into the `GRID_HOME\bin` folder. The drivers are copied into the `%systemroot%\system32\drivers` folder.

For Oracle Restart configurations, Oracle ACFS, Oracle Kernel Services Driver (OKS), and Oracle ADVM drivers are installed and initially loaded into the operating system kernel memory during execution of the Oracle grid infrastructure root script.

However, if the Oracle Restart software stack is restarted, the three drivers must be manually loaded into the operating system kernel memory by executing the driver load command `acfsload` as shown in [Example B-1](#). To execute the command, the user must have root privileges or Windows Administrator privileges.

Example B-1 Loading Drivers

```
grid_home/bin/acfsload start -s
```

File systems in the Oracle ACFS mount registry must also be manually remounted. You can mount all registered file systems in a Linux environment by executing the command shown in [Example B-2](#). A dummy name (`none`) was required for both the device and directory path even though these names are ignored with the `all` option. To execute the command, the user must have root privileges. For information about the `mount` command, see "[mount](#)" on page 13-10.

Example B-2 Mounting File Systems

```
# /bin/mount -t acfs -o all none none
```

The Oracle ASM instance is started during the Grid Infrastructure installation process whenever the Oracle Clusterware Registry (OCR) and voting disks are configured within an Oracle ASM disk group. In that case, the Oracle ACFS drivers are initially loaded during Grid Infrastructure Installation as a result of the resource dependency. The Oracle ASM instance can also be started using the Oracle ASM Configuration Assistant and the Oracle ACFS drivers are loaded as a result of that action. In steady state mode, the Oracle ACFS drivers are automatically loaded during Oracle Clusterware initialization when the Oracle High Availability Services Daemon (OHASD) calls the start action for the Oracle ASM instance resource that also results in loading the Oracle ACFS drivers due to the resource dependency relationship. The start action for the Oracle ACFS drivers resource attempts to load the Oracle ACFS, Oracle ADVM, and OKS drivers into the native operating system.

The policy for the Oracle ACFS drivers is that they remain loaded until the Oracle Clusterware is shut down. The `ora.drivers.acfs` resource is managed automatically by Oracle High Availability Services Daemon (OHASD) and its state cannot be manually manipulated by `srvctl` or `crsctl`.

Oracle ACFS Registry Resource Management

The Oracle ACFS registry resource is supported only for Oracle grid infrastructure cluster configurations; it is not supported for Oracle Restart configurations. See "[Oracle ACFS and Oracle Restart](#)" on page B-6.

The Oracle ACFS registry resource (`ora.registry.acfs`) is created by the root script that is executed following Grid Infrastructure installation. The start action for the Oracle ACFS mount registry resource is automatically called during Grid Infrastructure initialization to activate the local node state of the clusterwide Oracle ACFS mount registry. If this initialization is successful, the state of this resource is set to `online`; otherwise, the state of the resource is set to `offline`. The state of the Oracle ACFS registry resource is determined only by the active state of the mount registry. The `online` status is independent of any registry contents or the current state of any individual registered file systems that may exist within the Oracle ACFS registry.

In addition to activating the local node state of the mount registry, the Oracle ACFS registry resource start action assists in establishing a clusterwide Oracle ACFS file name space. On each node, the resource start action scans the contents of the clusterwide mount registry and mounts any file systems designated for mounting on the local cluster member. Before mounting a registered file system, the resource start action confirms that the associated file system storage stack is active and will mount the disk group, enable the volume file, and create the mount point if necessary to complete the mount operation.

The check action for the Oracle ACFS registry resource assists in maintaining the clusterwide Oracle ACFS file system name space. On each node, the check action scans

the contents of the mount registry for newly created entries and mounts any Oracle ACFS file systems registered for mounting on the local node. As a result a new Oracle ACFS file system can be created and registered on one node of the cluster, and is automatically mounted on all cluster members designated by the Oracle ACFS registry entry.

The Oracle ACFS registry resource check action also assists with file system recoveries. Recovering a file system from an offline state requires dismounting and remounting the file system. As the Oracle ACFS registry resource check action scans the mount registry searching for newly created file systems, it also checks for any offline file systems on the local node and if found attempts to dismount and remount each offline file system. If the remount is successful, the file system transitions from offline to fully active status.

The Oracle ACFS registry resource stop action is usually called during the Grid Infrastructure shutdown sequence of operations. To transition the registry resource to an offline state, all file systems on this cluster member that are configured with Oracle ADVM devices must be dismounted. A mounted file system maintains an open reference on its Oracle ADVM device special file and associated dynamic volume file that must be closed before the Oracle ASM instance can be shutdown normally. The registry resource stop action scans the operating system's internal mount table searching for any mounted file system that is configured with an Oracle ADVM device file. If any is found, the stop action attempts to dismount that file system. However, if there are open references resulting from applications or users of that file system, then the file system cannot be dismounted until these are closed. If the dismount operation fails, the process IDs of any processes holding an open reference on the file system are displayed and logged to enable the administrator to resolve the open references and dismount the file systems. The internal mount table entries can include registered and unregistered Oracle ACFS file systems, and other local file systems that were mounted on an Oracle ADVM device file.

The Oracle ACFS registry resource clean action is called implicitly if the resource stop action fails to transition the resource to the offline state. In that case, the registry resource clean action can be called to effectively force the resource offline. The registry resource clean action scans the operating system internal mount table searching for any file system that is mounted upon an Oracle ADVM device. If any is found, the resource clean action attempts to unmount the file system as in the resource stop action. However, if there are open references that prevent the file system from being dismounted, the clean resource action displays and logs the Process Identifiers of any process holding a reference, terminates the referencing processes, and then dismounts the file system. At the completion of the clean action, the registry resource is set to an offline state and other participants in the Grid Infrastructure shutdown sequence can now be stopped.

Whenever Oracle Clusterware is started on a cluster node, the Oracle ACFS startup operations for the node consult the cluster mount registry and attempt to mount all Oracle ACFS file systems that are registered for this node. Following each file system addition to the mount registry, the newly registered file system is automatically mounted on each node designated by the registry entry. If a registered file system is automatically mounted and is later dismounted, it is not automatically remounted until the system is rebooted or the Oracle Clusterware is restarted. It can be manually remounted using the `mount` command or Oracle Enterprise Manager.

The Oracle ACFS cluster mount registry action routines attempt to mount each Oracle ACFS file system on its registered mount point and creates the mount point if it does not exist. The registry action routines also mount any Oracle ASM disk groups and enable any Oracle ADVM volumes required to support the Oracle ACFS mount operation. In the event that a file system enters into an offline error state, the registry

action routines attempt to recover the file system and return it to an on-line state by dismounting and remounting the file system. For information about the offline error state, see "[About Oracle ACFS Integration with Oracle ASM](#)" on page 5-7.

Oracle ACFS Individual File System Resource Management

The Oracle ACFS individual file system resource is supported only for Oracle grid infrastructure cluster configurations; it is not supported for Oracle Restart configurations. See "[Oracle ACFS and Oracle Restart](#)" on page B-6.

Oracle ASM Configuration Assistant (ASMCA) facilitates the creation of Oracle ACFS individual file system resources (*ora.diskgroup.volume.acfs*). During database creation with the Database Configuration Assistant (DBCA), the individual file system resource is included in the dependency list of its associated disk group so that stopping the disk group also attempts to stop any dependent Oracle ACFS file systems.

An Oracle ACFS individual file system resource is typically created for use with application resource dependency lists. For example, if an Oracle ACFS file system is configured for use as an Oracle Database home, then a resource created for the file system can be included in the resource dependency list of the Oracle Database application. This dependency causes the file system and stack to be automatically mounted as a result of the start action of the database application.

An Oracle ACFS file system that is to be mounted from a dependency action should not be included in the Oracle ACFS mount registry.

The start action for an Oracle ACFS individual file system resource is to mount the file system. This individual file system resource action includes confirming that the associated file system storage stack is active and mounting the disk group, enabling the volume file, and creating the mount point if necessary to complete the mount operation. If the file system is successfully mounted, the state of the resource is set to *online*; otherwise, it is set to *offline*.

The check action for an individual file system resource verifies that the file system is mounted. It sets the state of the resource to *online* status if mounted, otherwise the status is set to *offline*.

The stop action for an Oracle ACFS individual file system resource attempts to dismount the file system. If the file system cannot be dismounted due to open references, the stop action displays and logs the process identifiers for any processes holding a reference.

Use of the `srvctl start` and `stop` actions to manage the individual file system resources maintains their correct resource state.

Oracle ACFS and Oracle Restart

Oracle Restart does not support root-based Oracle ACFS resources for this release. As a result, the following operations are not automatically performed:

- Loading Oracle ACFS drivers
- Mounting Oracle ACFS file systems listed in the Oracle ACFS mount registry
- Mounting resource-based Oracle ACFS database home file systems

The Oracle ACFS resources associated with these actions are not created for Oracle Restart configurations.

While Oracle ACFS resource management is fully supported for Oracle grid infrastructure configurations, the Oracle ACFS resource-based management actions must be replaced with alternative, sometimes manual, operations in Oracle Restart configurations.

See Also:

- "Grid Infrastructure Requires Manual Restart of Oracle ASM Cluster Drivers" in *Oracle Database Release Notes for Linux*
- "Oracle Automatic Storage Management Cluster Drivers Require Manual Restart" in *Oracle Database Release Notes for Microsoft Windows*.

Understanding Oracle ACFS I/O Failure Console Messages

Oracle ACFS logs information for I/O failures in the operating-specific system event log.

A console message has the following format:

```
[Oracle ACFS]: I/O failure (error_code) with device device_name during a operation_name op_type.
file_entry_num Starting offset: offset. Length of data transfer: io_length bytes.
Impact: acfs_type Object: object_type Oper.Context: operation_context
Snapshot?: yes_or_no AcfsObjectID: acfs_object_id . Internal ACFS Location: code_location.
```

The italicized variables in the console message syntax correspond to the following:

- I/O failure

The operating system specific error code, in Hex, seen by Oracle ACFS for a failed I/O. This may indicate a hardware problem, or it might indicate a failure to initiate the I/O for some other reason.
- Device

The device involved, usually the ADVM device file, but under some circumstances it might be a string indicating the device minor number
- Operation name

The kind of operation involved:
user data, metadata, or paging
- Operation type

The type of operation involved:
synch read, synch write, asynch read, or asynch write
- File entry number

The Oracle ACFS File entry number of the file system object involved, as a decimal number. The `acfsutil info fileid` tool can be used to find the corresponding file name.
- Offset

The disk offset of the I/O, as a decimal number.
- Length of I/O

The length of the I/O in bytes, as decimal number.

- File system object impacted

An indication that the file system object involved is either node-local, or is a resource accessed clusterwide. For example:

Node or Cluster
- Type of object impacted

A string indicating the kind of file system object involved, when possible. For example:

Unknown, User Dir., User Symlink, User File, Sys.Dir, Sys.File, or MetaData

 - Sys.Dir.

Oracle ACFS-administered directory within the visible namespace
 - sys.File

Oracle ACFS-administered file within the visible namespace
 - MetaData

Oracle ACFS-administered resources outside of the visible namespace
- Operational context

A higher-level view of what code context was issuing the I/O. This is for use by Oracle Support Services. For example:

Unknown, Read, Write, Grow, Shrink, Commit, or Recovery
- Snapshot

An indication of whether, if possible to determine, the data involved was from a Snapshot. For example:

Yes, No, or ?
- Object type of the file system

An internal identifier for the type of file system object. For use by Oracle Support Services.
- Location of the code

An internal identifier of the code location issuing this message. For use by Oracle Support Services.

The following is an example from `/var/adm/messages` in a Linux environment:

```
[Oracle ACFS]: I/O failure (0xc0000001) with device /dev/sdb during a metadata synch write .
Fenum Unknown. Starting offset: 67113984. Length of data transfer: 2560 bytes.
Impact: Node   Object: MetaData   Oper.Context: Write
Snapshot?: ?   AcfsObjectID: 8   . Internal ACFS Location: 5 .
```

Glossary

Oracle Automatic Storage Management Cluster

An Oracle Automatic Storage Management (Oracle ASM) cluster is a collection of interconnected nodes, each with an Oracle ASM instance, operating as a unified cluster with using Oracle Clusterware. An Oracle ASM cluster presents a shared pool of storage to one or more Oracle Databases that are also operating on the nodes. The databases can also be clustered using Oracle Real Application Clusters, but that is not a requirement. The disks that provide the shared storage pool must be accessible from all of the nodes in the Oracle ASM cluster.

ASMLIB

ASMLIB is an application programming interface (API) developed by Oracle to simplify the operating system-to-database interface and to exploit the capabilities of vendors' storage arrays on Linux-based systems.

Cluster File System (CFS)

A Cluster File System, or CFS, is a file system that is distributed across multiple computer systems. Generally, the computer systems share access to a disk connected through a [Storage Area Network \(SAN\)](#). The CFS component on each individual computer system coordinates access to the disks using a global communication mechanism.

Cluster Synchronization Services (CSS)

Cluster Synchronization Services (CSS) provide the foundation for coordinated, inter-process communication activities across a set of nodes. CSS services include group services, lock services, node information and cluster configuration services, and node monitoring services that detect nodes joining and leaving the cluster. OFS interacts with CSS to maintain a consistent view of cluster membership.

CSS Cluster

A CSS cluster is the cluster configuration defined by the CSS cluster membership services. CSS maintains the authoritative membership for an Oracle cluster and monitors nodes joining and leaving the cluster configuration. There is at most one CSS cluster defined for a given set of nodes.

disk group

An Oracle ASM disk group is a collection of disks that Oracle ASM manages as a unit. Within a disk group, Oracle ASM exposes a file system interface for Oracle Database files. The content of files that are stored in a disk group are evenly distributed, or striped, to eliminate hot spots and to provide uniform performance across the disks.

Oracle ASM files may also be optionally mirrored within a disk group. The performance of disks in a disk group is comparable to the performance of raw devices.

Direct Attached Storage (DAS)

Direct Attached Storage, or DAS, consists of storage devices that attach directly to host without an intervening network. DAS generally costs less than SAN attached storage, but also offers less flexibility and functionality. Examples of DAS include SATA disks, common to most desktop systems, and SCSI disks that are found on many servers.

extent

An extent, also called data extent, is the raw storage used to hold the contents of a file. Each extent consists of one or more allocation units (AU) on a specific disk.

extent map

An extent map is a list of extent pointers that point to all the data extents of a file. This is the raw storage for the file. Each extent pointer gives the disk and allocation unit of the data extent. For reliability each extent pointer also includes a check byte to ensure it has not been damaged. This is particularly important when using an in-memory copy of the extent map to direct file I/O to a particular disk location.

fibre channel

Fibre channel is a set of standards that define the interface between computer systems and peripherals. The fibre channel interface is a serial bus interface originally designed to supersede the SCSI standard. However, both the fibre channel and SCSI interfaces have evolved independently of each other and have benefited from the existence of the other. Fibre is spelled with **re** rather than an **er** to indicate a distinction from terms such as *fiber optics*. Both copper and fiber optics are commonly used as a media for fibre channel.

file system

A file system is a software component providing structured access to disks. File systems present objects, such as files, to application programs. Access to files is generally specified with a standard API defining operating system calls such as Open/Close and Read/Write that the application program uses for accessing files. File systems are usually provided as a component of an operating system, but may be provided as an independent software component.

file

A file is an object presented to an application program by a **file system**. A file is subdivided into blocks by the file system. A file system typically places what appears to the application program as consecutive blocks, into nonconsecutive locations on disks. The mapping of the file blocks to the locations on disks is kept in what is known as an **extent map**.

host bus adapter (HBA)

A host bus adapter (HBA), or host adapter, connects a host system to other network and storage devices, such as, eSATA, and SCSI devices.

Logical Unit Number (LUN)

In storage technology, the term LUN is often used to denote a disk presented to a computer system by a storage array. In SCSI terminology, a Logical Unit Number, or LUN, is a number representing a device address presented to a computer system.

mirroring

In storage systems, mirroring is a means for protecting the integrity of data by storing copies of data on multiple disks. If a disk fails, then a secondary copy of the data is available on a second or third disk. Generally, mirroring can be deployed in one of two ways. In the most common case, as with a RAID storage array, a fixed amount of physical storage space on a single disk is uniformly copied on multiple disks. Through software on the storage array, the physically mirrored space on multiple devices is presented to the host as a single disk. Another approach to mirroring is to copy fixed segments of a file onto multiple disks. Oracle ASM uses the second approach for maximum flexibility.

Network Attached Storage (NAS)

Network Attached Storage, or NAS, comprises a class of systems providing file services to host computers. A device driver in the operating system accesses data using protocols such as NFS or Common Internet File System (CIFS). From the application program's point of view, Network Attached Storage appears as a file system local to the host computer on which the application program resides.

Network File System (NFS)

A network file system, or NFS, is an industry standard providing Network Attached Storage on many platforms. NFS extends the local file system framework to provide device sharing to users on an IP network. The user view of an NFS is that a remote file system on a host appears as if it were locally mounted.

partition

Operating systems typically provide a means for splitting a disk into sections called partitions. To an application, a partition appears as an independent disk, even though that is not the case. The operating system provides a command for managing the size and locations of partitions on a disk. This command writes a partition map onto the disk in an area that is not visible to the application program.

Redundant Array of Inexpensive Disks (RAID)

Redundant Array of Inexpensive Disks, or RAID, is a means for combining multiple disks through software and presenting them to a host as a collection of apparently distinct disks. Mirroring and striping are two means used for combining the disks as a collection. RAID can be implemented as a software component of a computer system or as software operating inside of a storage array. RAID operates on the physical presentation of storage as opposed to Oracle ASM mirroring and striping at the logical file level.

Storage Area Network (SAN)

Storage Area Network, or SAN, is a specialized network designed for attaching storage devices, such as disk arrays and tape drives, to computer systems. The most common network infrastructure used in SANs is **fibre channel**. SANs provide great flexibility by allowing a common interconnect infrastructure in which any storage device can be connected to any computer system. Zoning and LUN masking are commonly used for access control for SANs. Zoning is a mechanism where the SAN infrastructure determines the access rights of storage devices to computer systems. LUN masking is a mechanism where the storage device determines the access rights of what computer systems is permitted to connect to the device.

Small Computer System Interface (SCSI)

Small Computer System Interface, or SCSI, is a standard defining the interface between computer systems and peripherals, most commonly storage devices. The

SCSI standard defines both the logical and physical aspects of the interface between computer systems and peripherals.

striping

In storage systems, striping is a means for spreading data across multiple disks as opposed to storing the data on a single disk. Striping is usually done to improve performance. Generally, striping can be deployed in one of two ways. In the most common case, as with a RAID storage array, a fixed amount of physical storage space that could have been stored on a single disk is uniformly striped across multiple disks. Through software on the storage array, the physically striped space on multiple devices is presented to the host as a single disk. Another approach to striping is to stripe fixed segments of a file across multiple disks connected to the host. Oracle ASM uses the second approach for maximum performance.

volume

In the storage realm, the meaning of volume has many related definitions. Volumes are often thought of as the virtual entity represented as a **Logical Unit Number (LUN)**. Volumes often are presented as an aggregation of pieces from several disks. A volume is managed by a software component called a **volume manager**.

volume manager

A volume manager is a software component that manages the mapping of the collection of the pieces of the disks into a **volume**.

A

- access control
 - disk groups, 4-38
- ACCESS_CONTROL.ENABLED
 - disk group attribute, 4-2, 4-40
- ACCESS_CONTROL.UMASK
 - disk group attribute, 4-2, 4-40
- acfschkdsk command, 13-13
- acfsdbg command, 13-22
- acfsdismount command, 13-15
- acfsformat command, 13-16
- acfsmountvol command, 13-17
- acfsutil detach command, 13-18
- acfsutil info command, 13-25
- acfsutil registry command, 13-27
- acfsutil rmfs command, 13-28
- acfsutil size command, 13-29
- acfsutil snap create command, 13-30
- acfsutil snap delete command, 13-32
- acfsutil tune command, 13-32
- active session history sampling
 - Oracle Automatic Storage Management, 3-21
- adding
 - disks, 1-10
 - disks to a disk group, 9-9
 - templates to a disk group, 7-17
- administering
 - disk groups, 4-1
- administration
 - Oracle ACFS, 5-8
- advmutil dismount command, 13-18
- advmutil list command, 13-19
- advmutil mount command, 13-20
- advmutil tune command, 13-33
- advmutil volinfo command, 13-34
- aliases
 - dropping from a disk group, 7-10
 - managing in Oracle Automatic Storage Management, 7-9
- allocation unit size
 - disk groups, 4-7
- allocation units
 - different sizes enabled by compatibility attributes, 4-36
- ALTER DISKGROUP command, 4-11
- archive log files
 - creating in Oracle Automatic Storage Management, A-4
- ASM file access control commands
 - ASMCMD, 12-50
- ASM SPFILE in a disk group
 - enabled by compatibility attributes, 4-36
- ASM_DISKGROUPS initialization parameter, 3-7
- ASM_DISKSTRING initialization parameter, 3-7
 - discovering disks, 1-10
- ASM_POWER_LIMIT initialization parameter, 3-8
- ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter, 3-8
- asmadmin group, 3-24
- ASMCA
 - see* Oracle ASM Configuration Assistant
- ASMCMD
 - about, 12-1
 - ASM file access control commands, 12-50
 - cd command, 12-19
 - chdg command, 12-28
 - chgrp command, 12-50
 - chkdg command, 12-30
 - chmod command, 12-51
 - chown command, 12-52
 - chtpl command, 12-47
 - cp command, 12-19
 - disk group management commands, 12-28
 - dropdg command, 12-31
 - dsget command, 12-10
 - dsset command, 12-10
 - du command, 12-21
 - exit command, 12-7
 - file management commands, 12-18
 - find command, 12-21
 - groups command, 12-53
 - grpmod command, 12-53
 - help command, 12-9
 - including directory in prompt, 12-8
 - instance management commands, 12-9
 - iostat command, 12-31
 - ls command, 12-22
 - lsattr command, 12-33
 - lsct command, 12-11
 - lsdg command, 12-34
 - lsdsk command, 12-36

- lsgrp command, 12-54
- lsod command, 12-38
- lsop command, 12-24
- lsop command, 12-12
- lspwusr command, 12-12
- lstmpl command, 12-48
- lsusr command, 12-55
- md_backup command, 12-38
- md_restore command, 12-39
- mkalias command, 12-25
- mkdg command, 12-40
- mkdir command, 12-25
- mkgrp command, 12-55
- mktml command, 12-49
- mkusr command, 12-56
- mount command, 12-42
- offline command, 12-43
- online command, 12-43
- orapwusr command, 12-13
- passwd command, 12-57
- preparing to run, 12-6
- pwd command, 12-26
- rebal command, 12-44
- remap command, 12-45
- rm command, 12-26
- rmlalias command, 12-27
- rmgrp command, 12-57
- rmtmpl command, 12-49
- rmusr command, 12-58
- running in a database home, 12-6
- running in an Oracle grid infrastructure home, 12-6
- setattr command, 12-46
- shutdown command, 12-14
- spbackup command, 12-14
- spcopy command, 12-15
- specifying the verbose mode, 12-8
- spget command, 12-16
- spmmove command, 12-16
- spset command, 12-17
- startup command, 12-18
- SYSASM privilege, 12-8
- template management commands, 12-47
- type of connection, 12-8
- umount command, 12-46
- volcreate command, 12-59
- voldelete command, 12-60
- voldisable command, 12-61
- volenable command, 12-62
- volinfo command, 12-62
- volresize command, 12-63
- volset command, 12-64
- volstat command, 12-65
- volume management commands, 12-58
- ASMCMD utility, 12-1
- asmdba group, 3-24
- asmoper group, 3-24
- ASMSNMP user
 - for monitoring the Oracle ASM instance, 3-23
- attributes

- disk group compatibility, 4-33
- DISK_REPAIR_TIME, 4-27
- for disk groups, 4-2, 4-5
- in templates for Oracle Automatic Storage Management, 7-15

- AU_SIZE
 - different sizes enabled by compatibility attributes, 4-36
 - disk group attribute, 1-7, 4-2, 4-5

B

- backing up
 - Oracle Automatic Storage Management files, 9-23

- backup and restore
 - Oracle ACFS, 5-7

C

- capacity
 - managing in disk groups, 4-21

- cd command
 - ASMCMD, 12-19

- CELL.SMART_SCAN_CAPABLE
 - disk group attribute, 4-2

- chdg command
 - ASMCMD, 12-28

- chgrp command
 - ASMCMD, 12-50

- chkdg command
 - ASMCMD, 12-30

- chmod command
 - ASMCMD, 12-51

- chown command
 - ASMCMD, 12-52

- chtml command
 - ASMCMD, 12-47

- Cluster File System (CFS), Glossary-1

- clusterware resource types
 - Kernel Services Driver (OKS), 5-9
 - Oracle ACFS, 5-9
 - Oracle ADVM, 5-9

- compatibility
 - disk groups, 4-32, 9-9, 9-15
 - Oracle ASM and database instances, 3-1
 - reverting disk groups to previous settings, 4-37

- COMPATIBLE initialization parameter, 4-33

- COMPATIBLE.ADVM
 - disk group attribute, 4-2, 4-5, 4-33

- COMPATIBLE.ASM
 - disk group attribute, 4-2, 4-5, 4-33

- COMPATIBLE.RDBMS
 - disk group attribute, 4-2, 4-5, 4-33

- configuration page
 - Oracle Automatic Storage Management, 9-3

- configuring
 - Oracle Automatic Storage Management, 9-3

- control file
 - creating in Oracle Automatic Storage Management, A-3

- cp command
 - ASMCMD, 12-19
- creating
 - database using Oracle Automatic Storage Management, A-1
 - disk groups, 4-3
 - disk groups in a new installation, 4-6

D

- data extent, Glossary-2
- DB_CACHE_SIZE initialization parameter, 3-9
- default configuration
 - Oracle ADVM, B-3
- DIAGNOSTIC_DEST initialization parameter, 3-9
- Direct Attached Storage (DAS), Glossary-2
- directories
 - managing disk group, 7-11
- discovering disks, 1-10, 4-18
- disk discovery
 - about, 4-19
 - improving time, 4-20
 - in Oracle Automatic Storage Management, 4-18
 - rules, 4-20
- disk failure
 - in Oracle Automatic Storage Management, 4-25
- disk group compatibility, 9-9, 9-15
 - considerations for replicated environments, 4-38
 - setting, 4-34
- disk group management commands
 - ASMCMD, 12-28
- disk groups, 3-15
 - access control, 4-38
 - ACCESS_CONTROL.ENABLED attribute, 4-40
 - ACCESS_CONTROL.UMASK attribute, 4-40
 - adding disks to, 9-9
 - adding templates to, 7-17
 - administering, 4-1
 - allocation unit size, 4-7
 - altering membership of, 4-11
 - and Oracle Cluster Registry (OCR), 4-10
 - and voting files, 4-10
 - attributes, 4-2, 4-5
 - checking internal consistency, 4-44
 - compatibility, 4-32
 - compatibility attributes, 4-33
 - creating, 4-3
 - creating for a new installation, 4-6
 - creating for Oracle Automatic Storage Management, 9-7
 - creating with SQL*Plus, 4-3
 - discovering disks, 4-19
 - dropping, 4-44
 - dropping disks from, 4-14, 9-11
 - dropping in Oracle Automatic Storage Management, 9-12
 - managing capacity in, 4-21
 - manually rebalancing, 4-17
 - monitoring disk space usage for Oracle Automatic Storage Management, 9-13

- mounting, 1-10
- mounting and dismounting, 4-42
- mounting with force, 4-43
- names, 4-3
- performance and scalability, 4-30
- preferred read, 3-8
- quorum failure groups, 4-10
- rebalancing, 1-10
- renaming, 4-45
- resizing disks in, 4-16
- reverting compatibility, 4-37
- sector size, 4-8
- understanding, 1-6
- undropping disks in, 4-17

- disk regions
 - and rebalance operations, 4-16
 - intelligent data placement, 4-15
 - Oracle Automatic Storage Management, 4-15
- DISK_REPAIR_TIME
 - disk group attribute, 4-2, 4-27
- disks
 - adding, 1-10
 - discovering, 1-10, 4-18, 4-19
 - dropping, 1-10
 - preparing for Oracle Automatic Storage Management, 2-1
 - understanding, 1-7
- driver model
 - Oracle ACFS, 5-6
- drivers resource
 - Oracle ACFS, B-3
- dropdgm command
 - ASMCMD, 12-31
- dropping
 - aliases from a disk group, 7-10
 - disk groups, 4-44
 - disks, 1-10
 - disks from a disk group, 4-14
 - disks from Oracle Automatic Storage Management
 - disk group, 9-11
 - files from a disk group, 7-10
 - Oracle Automatic Storage Management
 - templates, 7-18
- dsget command
 - ASMCMD, 12-10
- dsset command
 - ASMCMD, 12-10
- du command
 - ASMCMD, 12-21
- dynamic views, 6-1

E

- Enterprise Manager
 - see* Oracle Enterprise Manager
- Exadata storage
 - enabled by compatibility attributes, 4-36
- exit command
 - ASMCMD, 12-7
- extent, Glossary-2

extent map, Glossary-2
extents
 variable size, 1-7, 1-8, 1-9, 4-22, 4-36

F

failure groups
 about, 4-24
 and mirroring, 4-23
 preferred read, 4-29
 quorum type, 4-10
fast mirror resync, 4-27, 9-5, 9-6
 enabled by compatibility attributes, 4-36
fibre channel, Glossary-2
file, Glossary-2
file access and security
 Oracle ACFS, 5-8
file access control
 enabled by compatibility attributes, 4-36
 Oracle Automatic Storage Management, 4-38
file management commands
 ASMCMD, 12-18
file systems, Glossary-2
file templates
 understanding, 1-9
filenames
 Oracle Automatic Storage Management, 7-2
files
 templates, 1-9
find command
 ASMCMD, 12-21
fsck command, 13-7

G

grid infrastructure configuration
 Oracle ACFS, 5-8
grid infrastructure installation
 Oracle ACFS, 5-8
grid plug and play (GPNP)
 retrieving the location of the Oracle ASM
 SPFILE, 12-16
 setting the location of the ASM SPFILE, 12-17
groups
 OSASM, 3-14, 3-16, 3-25, 3-26, 3-27, 9-2, 9-7, 9-8,
 9-10
groups command
 ASMCMD, 12-53
grpmod command
 ASMCMD, 12-53

H

help command
 ASMCMD, 12-9
host bus adapter (HBA), Glossary-2

I

individual file system resource
 Oracle ACFS, B-6

initialization parameter file
 Oracle Automatic Storage Management, 3-3
initialization parameters
 COMPATIBLE, 4-33
 Oracle Automatic Storage Management, 3-2
instance management commands
 ASMCMD, 12-9
INSTANCE_TYPE initialization parameter, 3-9
intelligent data placement
 and rebalance operations, 4-16
 disk regions, 4-15
 enabled by compatibility attributes, 4-36
I/O failure console messages
 Oracle ACFS, B-7
iostat command
 ASMCMD, 12-31

L

LARGE_POOL_SIZE initialization parameter, 3-10
Logical Unit Number (LUN), Glossary-2
ls command
 ASMCMD, 12-22
lsattr command
 ASMCMD, 12-33
lsct command
 ASMCMD, 12-11
lsdg command
 ASMCMD, 12-34
lsdsk command
 ASMCMD, 12-36
lsgrp command
 ASMCMD, 12-54
lsod command
 ASMCMD, 12-38
lsolf command
 ASMCMD, 12-24
lsop command
 ASMCMD, 12-12
lspwusr command
 ASMCMD, 12-12
lstmpl command
 ASMCMD, 12-48
lsusr command
 ASMCMD, 12-55

M

managing
 Oracle Automatic Storage Management
 templates, 7-15
managing capacity in disk groups, 4-21
managing file access control
 Oracle Automatic Storage Management, 9-16
managing file access control page
 Oracle Automatic Storage Management, 9-16
managing files
 Oracle Automatic Storage Management, 9-19
managing files page
 Oracle Automatic Storage Management, 9-19

- managing templates
 - Oracle Automatic Storage Management, 9-20
- managing templates page
 - Oracle Automatic Storage Management, 9-20
- managing users
 - Oracle Automatic Storage Management, 9-6
- md_backup command
 - ASMCMD, 12-38
- md_restore command
 - ASMCMD, 12-39
- memory management
 - Oracle Automatic Storage Management, 3-5
- migrating a database to Oracle Automatic Storage Management, 3-27
- mirror
 - fast resync, 4-27, 9-5, 9-6
- mirroring
 - and failure groups, 4-23
 - storage systems, Glossary-3
 - with failure groups, 1-6
- mkalias command
 - ASMCMD, 12-25
- mkdgc command
 - ASMCMD, 12-40
- mkdir command
 - ASMCMD, 12-25
- mkfs command, 13-9
- mkgrp command
 - ASMCMD, 12-55
- mktpl command
 - ASMCMD, 12-49
- mkusr command
 - ASMCMD, 12-56
- modifying a file
 - disk regions, 4-16
- monitoring performance
 - Oracle Automatic Storage Management, 9-22
- mount command, 13-10
 - ASMCMD, 12-42
- mount model
 - Oracle ACFS, 5-6
- mount registry
 - Oracle ACFS, 5-6
- mounting
 - disk groups, 1-10
- mounting and dismounting disk groups, 4-42
- mounting at startup, 3-15
- mounting disk groups
 - using force, 4-43

N

- namespace
 - Oracle ACFS, 5-6
- naming
 - disk groups, 4-3
- Network Attached Storage (NAS), Glossary-3
- Network File System (NFS), Glossary-3
- NOWAIT keyword
 - in REBALANCE clause, 4-11

O

- OCR in disk groups
 - enabled by compatibility attributes, 4-36
- offline command
 - ASMCMD, 12-43
- online command
 - ASMCMD, 12-43
- operating system authentication
 - for Oracle Automatic Storage Management, 3-26
- Oracle ACFS
 - see* Oracle Automatic Storage Management Cluster File System
- Oracle ADVN
 - see* Oracle Automatic Storage Management Dynamic Volume Manager
- Oracle ASM
 - see* Oracle Automatic Storage Management
- Oracle ASM Configuration Assistant
 - about, 11-1
 - addDisk command, 11-20
 - commands, 11-15
 - configureASM command, 11-16
 - configureParameter command, 11-17
 - configuring an Oracle ACFS, 11-11
 - configuring an Oracle ASM volume, 11-9
 - configuring disk groups, 11-6
 - createACFS command, 11-21
 - createDiskGroup command, 11-18
 - createVolume command, 11-20
 - creating an Oracle ACFS, 11-12
 - creating an Oracle ASM instance, 11-2
 - creating an Oracle ASM volume, 11-10
 - creating disk groups, 11-7
 - deleteASM command, 11-18
 - mounting or dismounting an Oracle ACFS, 11-11
 - running in command-line mode, 11-14
 - running in silent mode, 11-14
 - starting, 11-1
 - updating Oracle ASM instance parameters, 11-2
 - upgradeASM command, 11-16
 - upgrading an Oracle ASM instance, 11-3
- Oracle Automatic Storage Management
 - accessing files with the XML DB virtual folder, 7-12
 - accessing the home page, 9-1, 10-1
 - active session history sampling, 3-21
 - administering, 3-1
 - aliases, 7-9
 - authentication, 3-22
 - command-line utility, 12-1
 - components, 1-2
 - configuration page, 9-3
 - configuring initialization parameters, 3-2
 - creating a database in, A-1
 - creating archive log files in, A-4
 - creating control file in, A-3
 - creating files in the database, 7-7
 - creating redo logs in, A-2
 - creating tablespaces in, A-2
 - discovering disks, 1-10, 4-18

- disk discovery, 4-18
 - disk discovery string, 1-10
 - disk failures in, 4-25
 - disk groups, 1-6
 - disk regions, 4-15
 - disks, 1-7
 - downgrading an instance, 3-19
 - failure groups, 4-24
 - file access control, 4-38
 - file templates, 1-9
 - filenames, 7-2
 - initialization parameter file, 3-3
 - instances, 1-2
 - integration with Oracle ACFS, 5-7
 - managing file access control page, 9-16
 - managing files page, 9-19
 - managing templates page, 9-20
 - memory management, 3-5
 - migrating a database to, 3-27
 - mirroring and failure groups, 1-6
 - monitoring disk group space usage, 9-13
 - operating system authentication for, 3-26
 - overview, 1-1
 - password file authentication for, 3-27
 - performance page, 9-22
 - preparing disks, 2-1
 - rolling upgrade, 3-21
 - shutting down an instance, 3-16
 - starting up an instance, 3-13
 - understanding, 1-2
 - upgrading an instance with Oracle ASM Configuration Assistant, 11-3
 - upgrading an instance with Oracle Universal Installer, 3-17
 - users page, 9-6
 - views, 6-1
 - with Oracle Enterprise Manager, 9-1, 10-1
 - XML DB virtual folder, 7-12
- Oracle Automatic Storage Management Cluster File System
- about, 5-4
 - administration, 5-8
 - and Oracle Database homes, 5-4
 - backup and restore, 5-7
 - basic steps to access a file system, 13-4
 - basic steps to create a file system, 13-1
 - basic steps to dismount and remove a volume and file system, 13-6
 - basic steps to manage snapshots, 13-5
 - basic steps to manage with command-line tools, 13-1
 - clusterware resource types, 5-9
 - command-line tools, 13-1
 - driver model, 5-6
 - drivers resource, B-3
 - file access and security, 5-8
 - grid infrastructure configuration, 5-8
 - grid infrastructure installation, 5-8
 - individual file system resource, B-6
 - integration with Oracle Automatic Storage Management, 5-7
 - I/O failure console messages, B-7
 - mount model, 5-6
 - mount registry, 5-6
 - namespace, 5-6
 - ora.diskgroup.volume.acfs file system resource, 5-9
 - ora.drivers.acfs drivers resource, 5-9
 - ora.registry.acfs registry resource, 5-9
 - overview, 5-1
 - registry resource, B-4
 - snapshots, 5-6
 - supported file types, 5-1
- Oracle Automatic Storage Management Dynamic Volume Manager
- default configuration, B-3
 - overview, 5-5, 5-10
- Oracle Automatic Storage Management File System
- views, 6-6
- Oracle Cluster Registry (OCR)
- in disk groups, 4-10
- Oracle Database homes
- and Oracle ACFS, 5-4
- Oracle Enterprise Manager
- support workbench for Oracle Automatic Storage Management, 9-25
- Oracle Restart
- about, 3-12
 - with Oracle Automatic Storage Management, 3-12
- ora.diskgroup.volume.acfs file system resource
- Oracle ACFS, 5-9
- ora.drivers.acfs drivers resource
- Oracle ACFS, 5-9
- orapwusr command
- ASMCMDB, 12-13
- ora.registry.acfs registry resource
- Oracle ACFS, 5-9
- OSASM group, 3-14, 3-16, 3-24, 3-25, 3-26, 3-27, 9-2, 9-7, 9-8, 9-10
- OSDBA group, 3-24
- OSOPER group, 3-24
- OSTMG10048, 4-18
- overview
- Oracle ACFS, 5-1
 - Oracle ADVM, 5-5, 5-10
 - Oracle ASM, 1-1

P

- partitions
 - disks, Glossary-3
- passwd command
 - ASMCMDB, 12-57
- password file authentication
 - for Oracle Automatic Storage Management, 3-27
- performance page
 - Oracle Automatic Storage Management, 9-22
- preferred read failure groups, 3-8, 4-29, 4-30
 - administering, 4-29

- preparing disks
 - for Oracle Automatic Storage Management, 2-1
- privileges
 - ASMCMD, 12-8
 - SYSASM, 3-14, 3-16, 3-25, 3-26, 3-27, 9-2, 9-7, 9-8, 9-10
- PROCESSES initialization parameter, 3-10
- pwd command
 - ASMCMD, 12-26

Q

- quorum failure group, 4-10

R

- read errors
 - recovery, 4-26
- rebal command
 - ASMCMD, 12-44
- rebalance
 - tuning, 4-18
- REBALANCE NOWAIT clause, 4-11
- REBALANCE WAIT clause, 4-11
- rebalancing
 - disk groups, 1-10
- rebalancing a disk group, 4-17
- recovery
 - read errors, 4-26
 - write errors, 4-26
- redo logs
 - creating in Oracle Automatic Storage Management, A-2
- REDUNDANCY_LOWERED column
 - in V\$ASM_FILE, 4-21
- Redundant Array of Inexpensive Disks (RAID), Glossary-3
- registry resource
 - Oracle ACFS, B-4
- remap command
 - ASMCMD, 12-45
- REMOTE_LOGIN_PASSWORDFILE initialization parameter, 3-10
- renamedg tool
 - renaming disk groups, 4-45
- renaming disk groups, 4-45
- replicated environments
 - disk group compatibility considerations, 4-38
- resizing disks in disk groups, 4-16
- restricting at startup, 3-15
- resync
 - fast mirror, 4-27, 9-5, 9-6
- reverting compatibility settings
 - disk groups, 4-37
- rm command
 - ASMCMD, 12-26
- rmlalias command
 - ASMCMD, 12-27
- rmgrp command
 - ASMCMD, 12-57

- rmtmpl command
 - ASMCMD, 12-49
- rmusr command
 - ASMCMD, 12-58
- rolling upgrade
 - Oracle Automatic Storage Management, 3-21
- running
 - ASMCMD, 12-6
 - ASMCMD in a database home, 12-6
 - ASMCMD in an Oracle grid infrastructure, 12-6

S

- sector size
 - disk groups, 4-8
 - enabled by compatibility attributes, 4-36
- SECTOR_SIZE
 - disk group attribute, 4-3
- setattr command
 - ASMCMD, 12-46
- SHARED_POOL_SIZE initialization parameter, 3-10
- shutdown command
 - ASMCMD, 12-14
- Small Computer System Interface (SCSI), Glossary-3
- snapshots
 - Oracle ACFS, 5-6
- spbackup command
 - ASMCMD, 12-14
 - backing up an Oracle ASM SPFILE, 12-14
- spcopy command
 - ASMCMD, 12-15
 - copying an Oracle ASM SPFILE, 12-15
- spget command
 - ASMCMD, 12-16
 - retrieving the location of an ASM SPFILE, 12-16
- spmmove
 - moving an Oracle ASM SPFILE, 12-16
- spmmove command
 - ASMCMD, 12-16
- spset command
 - ASMCMD, 12-17
 - setting the location of an ASM SPFILE, 12-17
- startup
 - mounting disk groups, 3-15
 - restricting disk groups, 3-15
- startup command
 - ASMCMD, 12-18
- Storage Area Network (SAN), Glossary-3
- striping
 - storage systems, Glossary-4
- supported file types
 - Oracle Automatic Storage Management Cluster File System, 5-1
- SYSASM privilege, 3-14, 3-16, 3-25, 3-26, 3-27, 9-2, 9-7, 9-8, 9-10

T

- tablespaces
 - creating in Oracle Automatic Storage

- Management, A-2
- template management commands
 - ASMCMD, 12-47
- templates
 - adding to a disk group, 7-17
 - attributes, 7-15
 - dropping, 7-18
 - for specifying redundancy and striping, 7-18
 - managing Oracle Automatic Storage Management, 7-15
 - modifying an Oracle Automatic Storage Management, 7-18
 - user-defined for Oracle Automatic Storage Management, 7-18

U

- umount command, 13-12
 - ASMCMD, 12-46
- undropping disks in disk groups, 4-17
- users page
 - Oracle Automatic Storage Management, 9-6

V

- V\$ASM_ACFSSNAPSHOTS view, 6-6
- V\$ASM_ACFSVOLUMES view, 6-6
- V\$ASM_ALIAS view, 6-1
- V\$ASM_ATTRIBUTE view, 4-36, 6-1
 - enabled by compatibility attributes, 4-36
- V\$ASM_CLIENT view, 6-1
- V\$ASM_DISK view, 6-1
- V\$ASM_DISK_IOSTAT view, 6-1
- V\$ASM_DISK_STAT view, 6-2
- V\$ASM_DISKGROUP view, 4-36, 6-2
- V\$ASM_DISKGROUP_STAT view, 6-2
- V\$ASM_FILE view, 6-2
- V\$ASM_FILESYSTEM view, 6-6
- V\$ASM_OPERATION view, 6-2
- V\$ASM_TEMPLATE view, 6-2
- V\$ASM_USER view, 6-2
- V\$ASM_USERGROUP view, 6-2
- V\$ASM_USERGROUP_MEMBER view, 6-2
- V\$ASM_VOLUME view, 6-6
- V\$ASM_VOLUME_STAT view, 6-6
- variable size extents, 1-7, 1-8, 1-9, 4-22
 - enabled by compatibility attributes, 4-36
- views
 - displaying information, 6-1
- volcreate command
 - ASMCMD, 12-59
- voldelete command
 - ASMCMD, 12-60
- voldisable command
 - ASMCMD, 12-61
- volenable command
 - ASMCMD, 12-62
- volinfo command
 - ASMCMD, 12-62
- volresize command

- ASMCMD, 12-63
- volset command
 - ASMCMD, 12-64
- volstat command
 - ASMCMD, 12-65
- volume management commands
 - ASMCMD, 12-58
- volume manager, Glossary-4
- volumes
 - enabled by compatibility attributes, 4-36
 - storage systems, Glossary-4
- voting files
 - in disk groups, 4-10
- voting files in disk groups
 - enabled by compatibility attributes, 4-36

W

- WAIT keyword
 - in REBALANCE clause, 4-11
- write errors
 - recovery, 4-26

X

- XML DB
 - virtual folder for Oracle Automatic Storage Management, 7-12